FERMILAB-Conf-96/069

Volume 2

Proceedings of the 1995 International Conference on Accelerator and Large Experimental Physics Control Systems

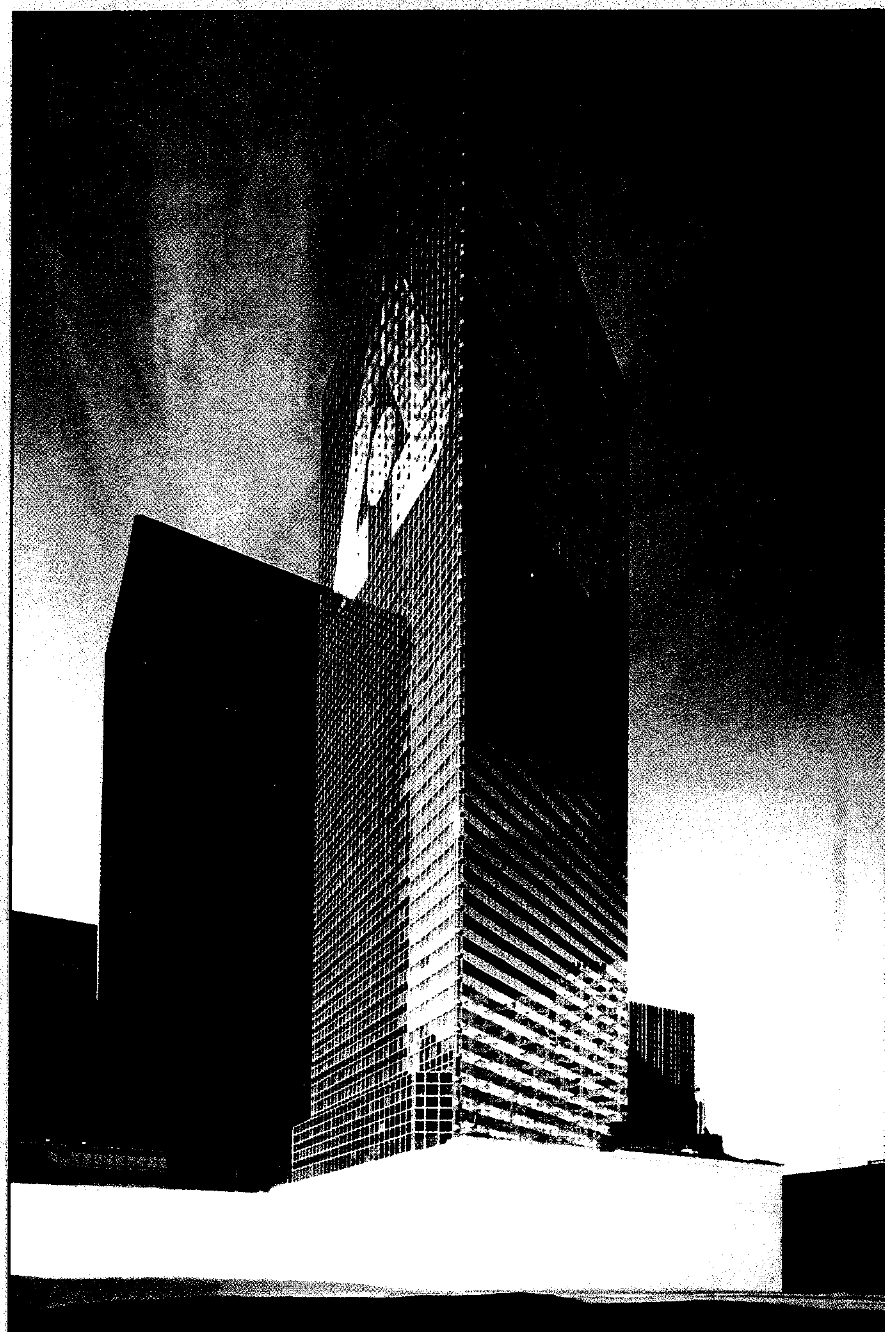

Chicago, Illinois

October 30 - November 3, 1995

FERMILAB Conf-96/069

Proceedings of the 1995 International Conference on Accelerator and Large Experimental Physics Control Systems

Volume 2

Chicago, Illinois
October 30-November 3, 1995

Editors

M.C. Crowley-Milling
St. Cergue, Switzerland

P. Lucas
Fermi National Accelerator Laboratory
Batavia, Illinois

P. Schoessow
Argonne National Laboratory
Argonne, Illinois

The DØ High Voltage System

S.Ahn, F.Bartlett, S.Fuess, S.Krzywdzinski, B.MacKinnon, L.Paterno
Fermi National Accelerator Laboratory
Batavia, IL 60510, USA

ABSTRACT

The DØ experiment at the Fermilab Tevatron Collider uses approximately 1000 independently controlled high voltage supplies in the detector. The VME-based front-end of the high voltage system employs Motorola 68020 processors running the pSOS operating system. Host processes, running on VAX/VMS computers, communicate to the front-end computers via a local area network. The DØ Control Data Acquisition package, CDAQ, handles all the communication between the host and the front-end. The front-end hardware and software is structured to ensure robustness and flexibility. The host software provides intelligent and intuitive high-level control and monitoring of the entire DØ high voltage system, which helps optimize the response of detector components.

I. OVERVIEW OF THE SYSTEM

The DØ [1] detector consists of several components which have distinct high voltage requirements. Within a detector component high voltage modules of differing characteristics are used. The design goal of the DØ high voltage system was to optimize the detector performance by meeting the individual and collective high voltage needs for the 1000 high voltage supplies. The operation of the detector also demands an uninterrupted and accurate supply of high voltage for long periods of time. Reliability is the most critical requirement for the system. The data taking conditions change frequently and requires the control and monitoring of the entire system with minimum effort. The high voltage system must be robust, flexible, intelligent and intuitive. The DØ high voltage system meets most of these requirements.

The high voltage system consists of front-end systems and host-level processes which communicate through the standard DØ control path [2]. An individual front-end is a VMEbus-based system with six Motorola 68020 microprocessors handling up to 192 high voltage supplies. DØ employs 7 functionally identical replications of this front-end unit. The host contains a set of processes on VAX/VMS which are part of the DØ control and monitoring system. The host processes use CDAQ services to communicate with front-end nodes which reside on token-ring.

In the following sections we will discuss the various components of the system and our experience with it.

II. HARDWARE

The DØ high voltage hardware consists of VMEbus-based high voltage modules [3], a crate controller, micro-processors that control the high voltage hardware and a token ring for communication with external systems.

A. High Voltage modules

A high voltage module consists of a control board and 8 independently controllable high voltage supplies. There are 7 types of supplies with differing resolutions and voltage and current limits. The resolution is determined by the range of the 15 bit ADC and the maximum voltage limit. A typical supply provides up to 5.5 kV at a current of 1 mA.

High voltage is generated by a pulse-width modulator and a Cockroft-Walton voltage multiplier (fig.1) The supplies deliver voltage linearly over a wide range. The accuracy of the delivered voltage is largely determined by the noise level of the high voltage components which are typically 0.05% of the full voltage. The supplies also exhibit long term stability when properly installed and maintained. For the 5.5 kV supply the long term voltage stability is better than 10 volts over 6 months.

The system high voltage monitor periodically reads high voltage data to determine the status of each supply. The status is represented using color block characters. The high voltage status of the entire DØ and a brief summary are shown in the display (fig.4). The arrangement of the high voltage channels closely resembles their physical arrangement. The status condition for each channel is defined by algebraic and logical relations. The definition is compiled to an interpretive form which is executed at run time. The definitions can also be inherited from another channel. All status changes are recorded in a database. The system high voltage monitor offers many tools that can be used in conjunction with regular periodic activities, but in practice user commands are rare.

The user interface to all the host processes was designed using the VAX Screen Management facility(SMG). It was originally designed to work in environments which did not have X display capabilities.

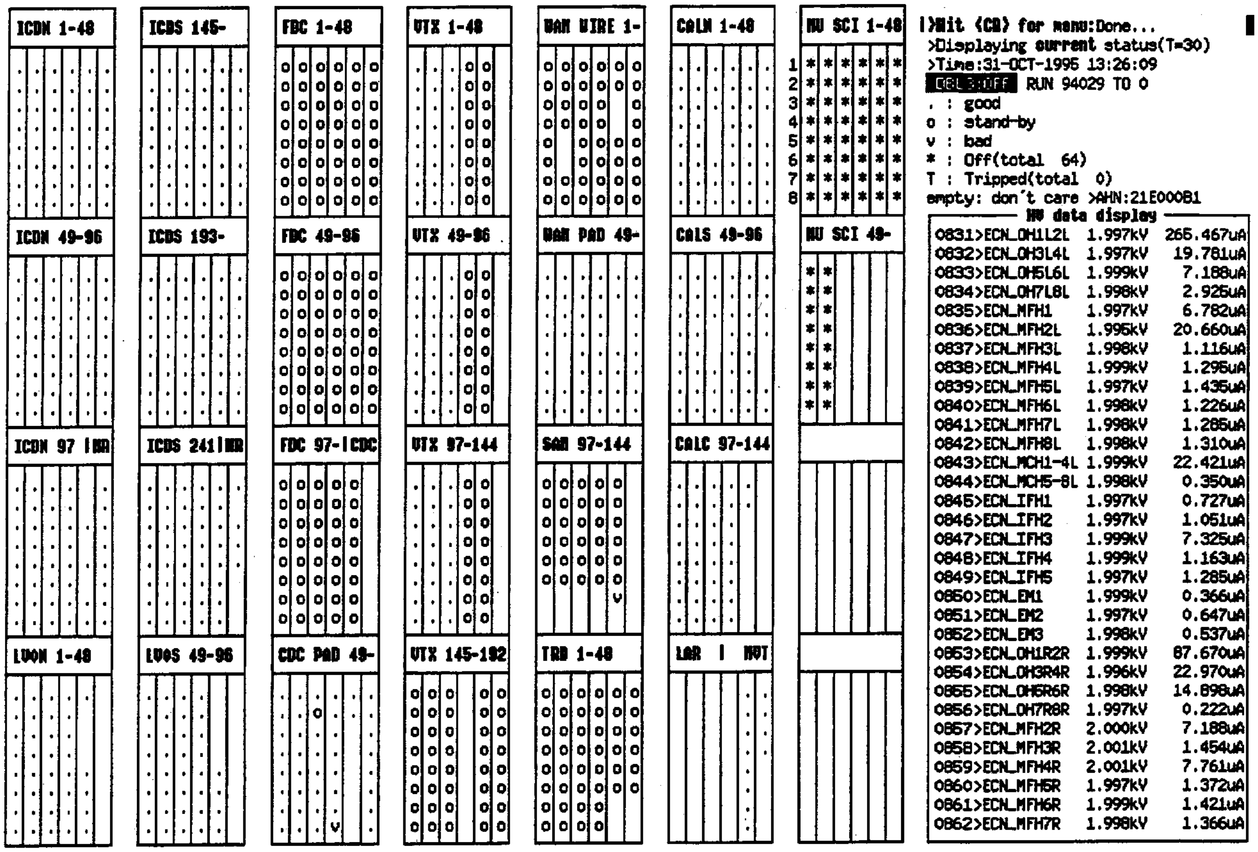

Fig.4 HVMON display

C. *Diagnostics Support Package*

A number of diagnotics tools were created for the high voltage system. The command queue buffer was used to debug command execution problems. About 1.5 days of slow history data is kept on the front-end to study the long term high voltage behavior of the detector. The fast history is used to debug both the detector and the high voltage system. The history data is accessible from the control process, the system high voltage monitor and other stand-alone processes. The history data is stored in Ntuple files, which can be viewed using PAW.

IV. EXPERIENCES AND LESSONS

Initially, DØ developed PC applications for controlling the high voltage. The PC environment was particularly useful during the hardware development phase and is now occasionally used when there are network disruptions which prevent the use of the operational system. It can also calibrate high voltage electronics and has an extensive help facility. However, because it was rarely used during the running of the experiment, casual users did not remember the command sequences and felt uncomfortable with it. We think an identical user interface for both the PC and host system would have avoided this problem.

ACKNOWLEDGMENTS

The authors wish to acknowledge the considerable contributions of the members of the Accelerator and Computing Divisions of the Fermi National Accelerator Laboratory as well as those of the DØ Collaboration.

REFERENCES

[1] S. Abachi, et al., (Fermilab PUB93-179-E) Nucl. Instr. and Meth. **A338** (1994) 185.
[2] S. Ahn, et al., (Fermilab PUB93-324-E) Nucl. Instr. and Meth. **A352** (1994) 250-253.
[3] S. Ahn, et al., Conference Record of the 1991 IEEE Nuclear Science Symposium. Vol *2*, p984.

MODEL BASED, DETAILED FAULT ANALYSIS IN THE CERN PS COMPLEX EQUIPMENT

M. Beharrell, G. Benincasa, J.M. Bouché, J. Cuperus, M. Lelaizant, L. Merard
PS Division, CERN, CH-1211 Geneva 23, Switzerland

Abstract

In the CERN PS Complex of accelerators, about a thousand pieces of equipment of various types (power converters, RF cavities, beam measurement devices, vacuum systems, etc.) are controlled using the so-called Control Protocol. This Protocol, a model-based equipment access standard, provides, amongst other facilities, a uniform and structured fault description and report feature. The faults are organized into categories according their severities and are presented at two levels, the first being global, and identical for all devices, and the second being very detailed and adapted to the peculiarities of each device.

All the relevant information is provided by the equipment specialists and is appropriately stored in static and real time databases; in this way a unique set of data-driven application programs can always cope with existing and newly added equipment.

Two classes of applications have been implemented, the first one is intended for control room alarm purposes and the second is oriented for specialist diagnostics. The system is completed by a fault history report facility permitting easy retrieval of faults which have occurred previously e.g. during the night.

INTRODUCTION

The control system of the CERN PS Complex, with its nine separate accelerators, deals with thousands of pieces of equipment of various types and having different control sequences from each other.

The maintenance of such a quantity of hardware requires the use of an appropriate set of tools that, on the one hand provides a detailed specific fault report system, possibly organized in a hierarchical structure, and on the other provides uniformity in the presentation of information. The Mean Time To Repair (MTTR) a device is in fact strongly dependent on the clarity and unambiguity of the information (fault messages) that is sent by the device and that must be interpreted by the maintenance team.

A particular aspect in the fault recovery treatment concerns intermittent faults, i.e. those usually having non-destructive effects and that occur at certain moments and disappear before appropriate treatment can be undertaken. The cumulative effects of these faults can be dangerous because they often precede some major problem in the device; the fault hunting system must then provide adequate tools for detecting and reporting these.

Finally, different faults occurring in the same device or in separate devices can be connected with each other or have the same root cause The detection of the relationship between different faults is greatly helped by appropriately flagging fault messages with time information.

All these considerations have been taken into account in the design of the fault report system presented in this paper.

1. THE PS COMPLEX CONTROL ENVIRONMENT

The nine accelerators of the PS Complex are controlled using the same unique architecture [1]. The control system is based on the so-called "standard model" for controls, i.e. an architecture having two (or three) levels of computing, interconnected by an appropriate LAN.

The first level is composed of powerful workstations running under UNIX and connected with each other and with the second level through an Ethernet LAN using TCP/IP. At the second level we find a series of VME crates (called Device Stub Controllers, DSC) housing 32-bit processors of the 680xx type and running LYNXOS RT. For certain kinds of equipment a third level of control is used, based on field buses.

Software access to the equipment is performed through standard control modules called Equipment Modules (EM) [2] ; these modules, one for each type of equipment, are housed in the DSCs, and hide to application programs all the intricacies of the control system. A large ORACLE database is housed in a dedicated server and contains all the necessary information to run the accelerators [3].

In order to improve the access speed at run-time, two subsets of this database are extracted:

- Data Base Real Time (DBRT) is housed in a local server (one per accelerator) and contains, amongst other information, the addresses of all controlled pieces of equipment
- Data Table (DT) constitutes an essential part of the EMs and is thus housed in the DSCs and contains all the information necessary to control each single piece of equipment physically connected to that DSC.

The Alarm System [4] periodically scans (polls) all the equipment of the Complex and reports to the workstations the faulty situations. In this context has been installed some years ago the so-called Control Protocol [5, 6], a model based uniform access procedure for equipment. All the equipment of the PS Complex has been classed in families, each family containing the devices having similar goals and characteristics (power converters, RF Cavities, beam diagnostic instrumentation, vacuum systems, etc.). For each family behavioral static and dynamic models have been defined and the corresponding control and acquisition parameters have been identified.

The implementation of the Control Protocol is based on two separate software packages exchanging appropriate control and acquisition messages. The first, one package for each family of devices, is hardware independent and is realized in the form of an Equipment Module permitting access to all the equipment of the same family. The second, one package for each single device, is strongly hardware dependent and implements all the control sequences peculiar to this device. The two packages are largely independent each other and can be written by the most appropriate persons; only the exchanged control and acquisition messages must conform with the Protocol rules, i.e. contain the parameters identified in the concerned equipment model, expressed in an appropriate formalism.

In the equipment models and, as a consequence, in the contents of exchanged messages the identification and report of fault conditions has received particular care.

2. THE FAULT DESCRIPTION IN THE CONTROL PROTOCOL

In the Control Protocol the faults or, better, the anomalous situations, are described at two levels:

- the level, called QUALIFIER, is a global description common and identical for all families
- the second , called DETAILED STATUS, is specific for each piece of equipment

The QUALIFIER contains six indicators, not exclusive each other, in increasing order of severity:

- WARNING indicates a minor fault not having consequences on the behavior of the device. Nevertheless, it may indicate the beginning of some more serious condition.
- BUSY only indicates that the device is executing some time-consuming task, e.g. a motor is still running or some sub-piece is warming-up to the correct temperature, etc. The Busy condition is accompanied, where possible, by indication of the number of seconds necessary to terminate the action.
- RESETTABLE FAULT indicates that a major fault has occurred and prevents normal behavior of the device; the recovery from this fault can usually be obtained by executing a computer controlled reset that tries to restart the hardware and software of the device.
- NON RESETTABLE FAULT indicates that a major fault has occurred, similar to the previous one, but that no recovery action by computer is possible and that the specialist must be called.
- INTERLOCK indicates that a fault occurred not directly in the concerned device but in an interconnected system, and that this fault prevents the normal behavior of the device. An example is a bad vacuum condition which can prevent the measurements in a beam diagnostic device.
- INTERNAL COMMUNICATION fault (not present for all devices) concerns those devices having distributed hardware interconnected with a local network.

These Fault indicators provided by the Qualifier often represent sufficient information on the status of the device for the operators in Control Room. On the contrary they do not contain enough information for the device specialist or for a more precise diagnostic. In fact each indicator is the sum of several possible faults of the same category: for example, the Resettable Fault indicates that one or more resettable faults occurred, without indicating which ones. The DETAILED STATUS indicators answer this need.

Five of the fault indicators of the Qualifier (the Busy indicator has not been considered) have been detailed in order to contain up to 32 different causes of a fault. The number 32 is not magic, it is just sufficient in the PS control system environment. The 5 x 32 Detailed Status indicators are specific to each piece of equipment, and in general their meaning is different from one device to another. The specialist on each device defines the meaning of the 32 indicators for each category of fault, organizes them in increasing order of severity and provides appropriate error messages for each. A check is done by the controls specialists in order to avoid too cryptic messages or ambiguities, e.g. two messages referring to the same fault or two faults having the same message. Each of the fault categories is flagged with a timestamp indicating the time of occurrence of the most severe of the reported faults.

The Qualifier information is implemented in the PS control system as six (five for certain families) bits of a word contained in every acquisition message from a device. In this way, at each acquisition one has global information on the status of that device. The Detailed Status information is obtained with a specific request and comes as a special acquisition message containing, amongst other information, five (or four for certain families), 32-bits words, each bit representing a particular fault, Fig. 1.

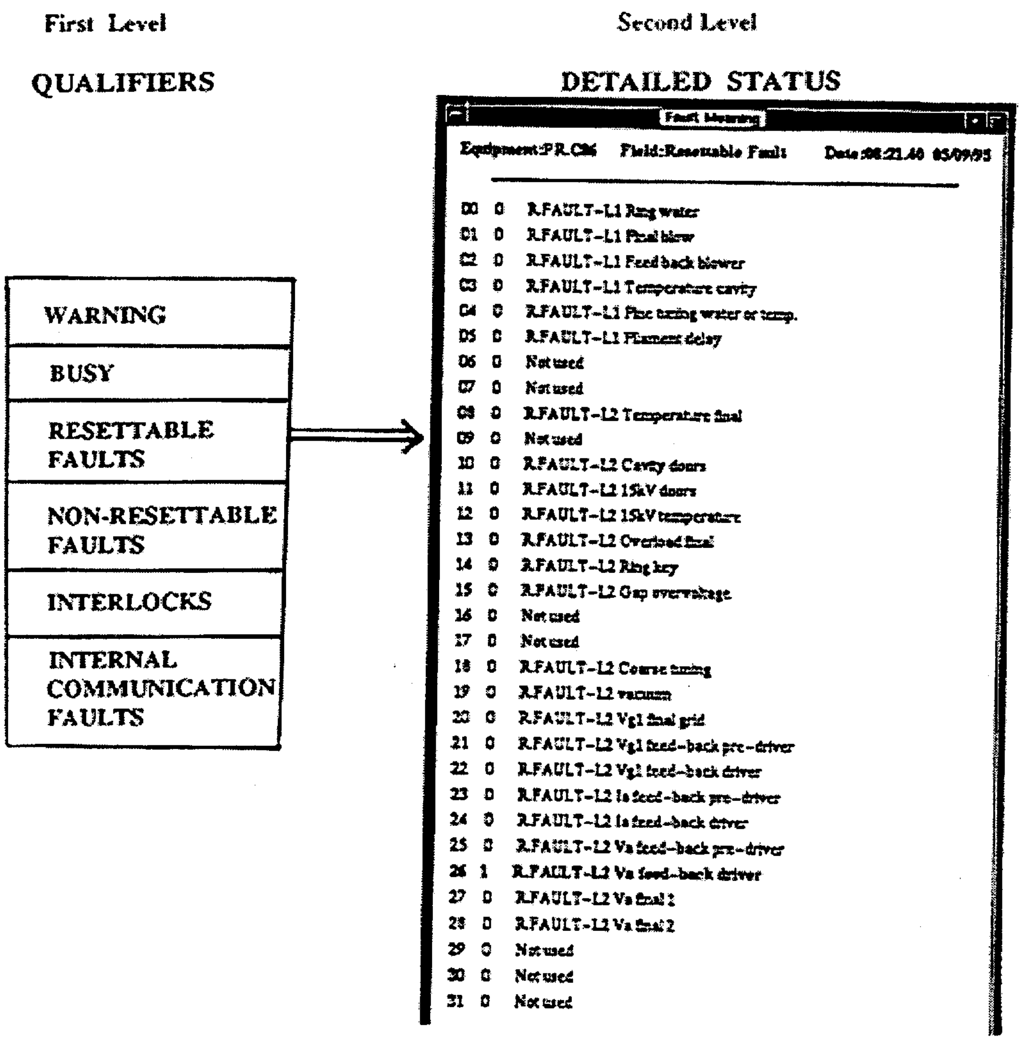

Fig.1 The two levels of the Fault report system (ex. a RF Cavity)

3. DATA BASE CLASSIFICATION OF THE FAULT INFORMATION

All the necessary information for the Fault report system is contained in two separate databases. A complete list of fault messages for all the PS Complex equipment is housed in the DBRT. The messages are stored and numbered in progressive order as they are provided by the equipment specialists; adding new messages is then a simple operation. Each message starts with indication of the category (Warning, Resettable fault etc..).

Appropriate functions exist that, given the fault number, retrieve the corresponding message. To date, about 400 fault messages are stored. The information specific to each device is housed in the corresponding Data Table in the concerned DSC. This information is stored in the form of 5 (4) groups of 32 integers, one group for each one of the previously mentioned categories of faults. The 32 integers in each group correspond to the 32 possible fault messages defined by the specialist for this specific category and for this device; each integer contains the number of fault message stored in the DBRT, as mentioned.

4. THE FIRST LEVEL OF APPLICATION - THE ALARM SYSTEM

As mentioned, the Alarm System periodically polls the various devices of the PS Complex and reports the anomalous situations in the form of short messages on the consoles. In the case where the polled device belongs to that class accessed by the Control Protocol, first of all the global fault descriptor Qualifier is acquired.

If one or more of the fault indicators are lighted, the function "Firstfault" is called. This function:

i) scans the Qualifier descriptor to identify, amongst the fault indicators, the most important in order of severity,

ii) interrogates the concerned device to obtain a Detailed Status Message, such as previously described. Among the 32 bits of fault in the selected category of indicator, it identifies the most important,

iii) finds the error message number corresponding to the position of the identified bit in the Data Table of the concerned device,

iv) retrieves, using this number as an input parameter the corresponding Fault Message from the DBRT and displays it on the console.

In this way the most important fault, amongst the 5 (4) x 32 possible fault conditions in this device, is displayed to the operator by the Alarm System.

5. THE SECOND LEVEL OF APPLICATIONS - THE DETAILED STATUS ACQUISITION

After a fault has been signaled by the Alarm System, somebody, either the specialist or the maintenance person, needs more information about the status of the concerned device. In this case the Detailed Status program is requested. This application is totally data driven in the sense that all the necessary information is contained in the DBRT and in the Data Tables - no new code or compilation is necessary when equipment is added or deleted in the accelerator complex.

In the example of Fig. 2 the case of a beam current transformer (TRAFO) is reported:

i) the application first searches in the DBRT to identify all the equipment families in the concerned accelerator using the Control Protocol facility; the complete list is then reported on the display,

ii) after the user has selected a family, the application searches once again in the DBRT for all the VME crates (DSC) containing the concerned kind of device and displays the list,

iii) the user clicks on the name of the DSC containing the hardware of the faulty device. By using the DBRT-stored information, the names (in the example beam transformers) of the concerned devices are displayed on the screen,

iv) finally the selection of the name of the device causes the request of a Detailed Status Message from the concerned hardware. This message contains (in this case) the 4x32 bits of fault information, as previously described. At this point the experienced specialist recognizes the various faults present in the device, often just by looking at the position of the set bits in the message,

v) otherwise, he can click on one group of 32 bits to obtain the messages on the display. To do this the application uses the information contained in the DBRT and in the Data Table, as already explained for the Alarm System. The complete list of the 32 fault messages is displayed, where the messages corresponding to actual faults are displayed in bold characters.

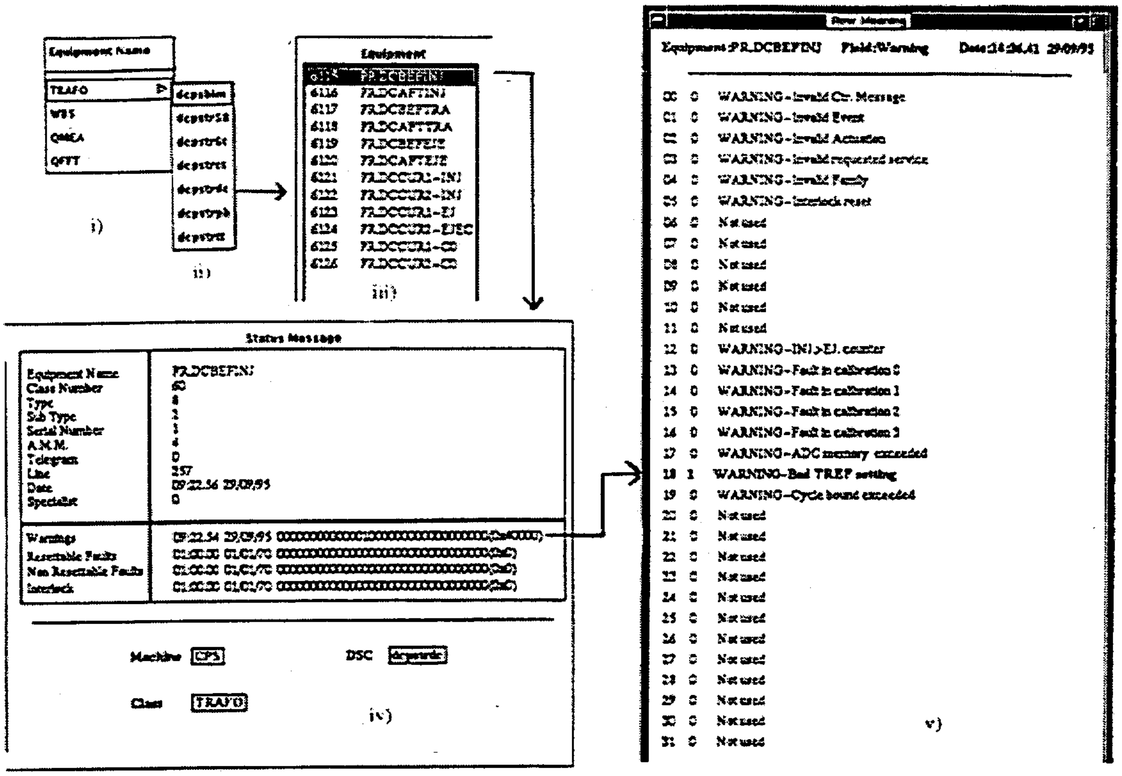

Fig. 2 The sequence in the Detailed Status application

6. THE FAULT HISTORY APPLICATION

As mentioned in the Section 1, certain faults are volatile and others are resettable by the operator before the specialist for the concerned device can be informed of their occurrence. On the other hand, it could be very useful to keep track of such faults for subsequent investigations. For this reason the Fault History system has been created.

The system is composed of two separate entities: a real time task, and an appropriate application. For various reasons, not reported here, almost all the devices of the PS Complex have a real time task, running in the corresponding DSC, and executing, among other activities, a periodical (~ 1.2 sec) acquisition of the relevant parameters which are subsequently stored in the Data Table.

For the devices using the Control Protocol facility, a complete acquisition message is acquired. In this case, when the specific software of a given device recognizes the occurrence of a fault that needs to be recorded, it sets an appropriate "look at me" flag in the next outgoing acquisition message. The real time task, after recognition of this flag, issues a request for a Detailed Status Message to the emitting device. The message with a timestamp is subsequently stored in an appropriate area of the Data Table organized in the form of a FIFO ring buffer. One such buffer exists per family of devices and per DSC:; it is presently sized to contain 100 Detailed Status Messages. In this way we have the complete record of the last 100 fault situations which have occurred, with their times.

The application program starts in the same way as described in i) and ii) of the previous section. At this point, by selecting the appropriate DSC name the program interrogates the ring buffer for this DSC and for the concerned family of devices, and displays the list of the recorded faulty elements with time, in chronological order. From this point on the program behaves in the same way as for iv) and v) of the previous chapter, the only difference being that now the Detailed Status messages are extracted from the ring buffer and not requested directly of the device.

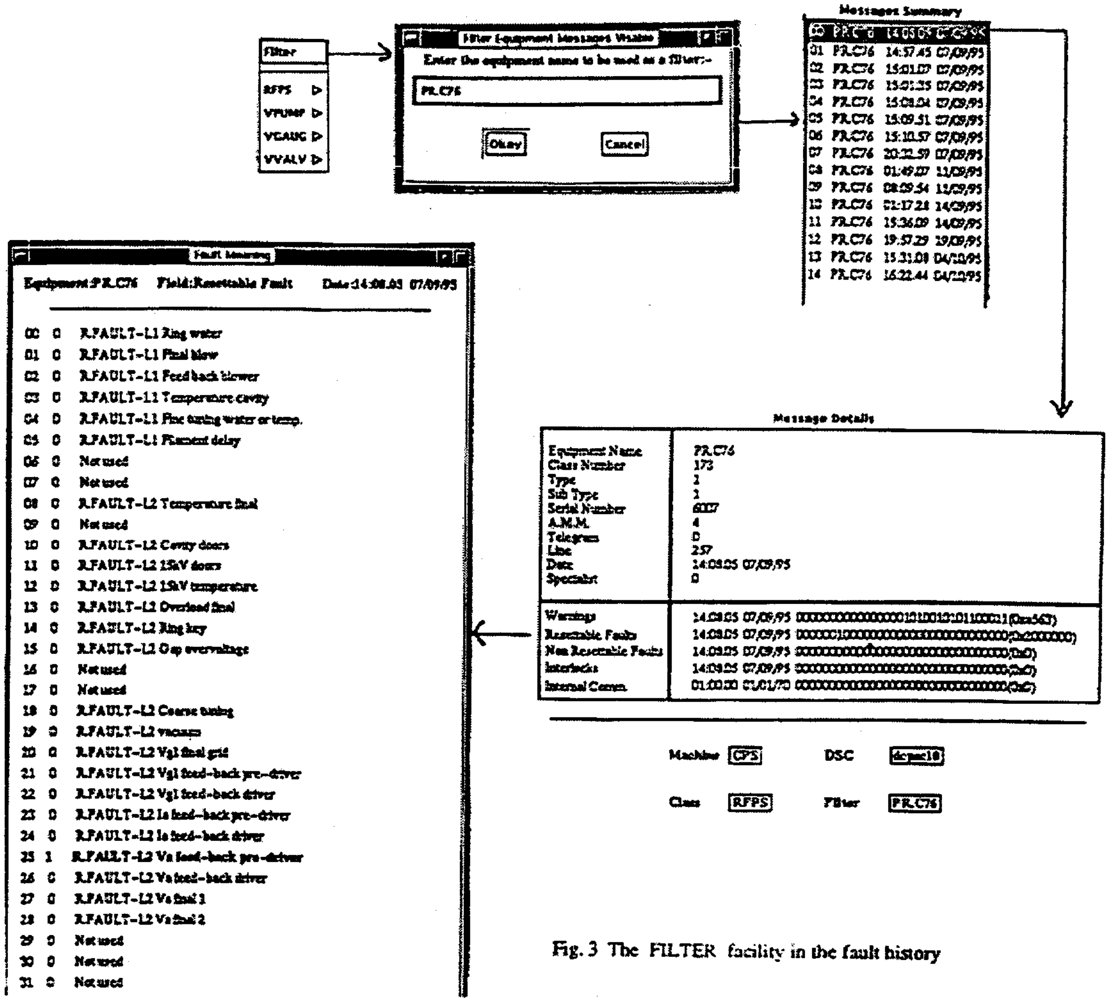

Fig. 3 The FILTER facility in the fault history

A useful feature, called Filter, permits easy following in time the fault behavior of a single device; to do this the user has only to provide the name of the inquired device, Fig. 3. The program searches the ring buffer and presents to the user a list of all occurrences of fault situations for the same device, chronologically organized. By examining the successive Detailed Status messages the specialist can extract important information on the fault's evolution.

7. CONCLUSIONS

In the PS Complex there are to date about a thousand pieces of equipment and devices using the Control Protocol; of this quantity, about half are thus far equipped with Detailed Status reports.

The presented standardization in Fault Reporting has proved to be a good compromise between the application of rules and constraints needed by any standard, and the necessity of implementing it in the accelerator environment where the variety of devices has in effect no limits.

The constraints have been limited only to the classification of the faults and in the formalism of their representation. This is especially appreciated in the control room where the operators receive information from thousands of devices. On the other hand, the equipment specialist is free to define the faults to be reported, the order of their severity, the corresponding fault messages, and that fault that should be recorded in the history.

REFERENCES

[1] F. Perriollat, C. Serre, " The new CERN PS control system; overview and status", ICALEPCS '93 , Berlin, Germany, October 18-23, 1993, Nucl. Instr. And Meth. A352 (1994) 86.
[2] J. Cuperus, W.Heinze, C.-H. Sicard, " Control Module Handbook", Version 4, CERN Internal Note PS/CO 95-01, 1995.
[3] J. Cuperus, " The database for accelerator control in the CERN PS Complex", IEEE Conference on Particle Accelerators, Washington, D.C. , March 1987.
[4] J. M. Bouché', J. Cuperus, M. Lelaizant , " The data driven Alarm System for the CERN PS accelerator Complex", ICALEPCS '93, Berlin, Germany, October 18-23, 1993, Nucl. Instr. And Meth. A352 (1994) 196.
[5] Reported by G. Benincasa, (USAP Members: G. Benincasa, O. Berrig, P. Burla, A. Burns, P. Charrue, G. Gelato, R. Keyser, K.-H Kissler, T. Linnearc, F. Perriollat, J. Pett, M. Ranany, C.-H. Sicard, P. Strubin), "Final report on the uniform equipment access at CERN", CERN Internal Report PS 93-16 (CO), 1993.
[6] G. Baribaud, I. Barnett, G. Benincasa, O. Berrig, R. Brun. P. Burla, R. Cappi, G. Coudert, C. Dehavay, B. Desforges, R. Gavaggio, G. Gelato, H.K. Kuhn, J. Pett, R. Pittin, J.P. Royer, E. Schulte, C. Steinbach, P. Strubin, D. Swoboda, N. Trofimov " Control Protocol : the proposed new CERN Standard access procedure to accelerator equipment - Status Report", ICALEPCS '91, Tsukuba, Japan, KEK Proc. 92-15 (1992) 591.

The Use of ARTEMIS With High-Level Applications*

B. A. Bowling, H. Shoaee, S. Witherspoon, J. Van Zeijts

Continuous Electron Beam Accelerator Facility, Newport News, VA 23606 USA

Abstract

ARTEMIS is an on-line accelerator modeling server developed at CEBAF. One of the design goals of ARTEMIS is to provide an integrated modeling environment for high-level accelerator control codes, for example, an automated beam steering application. This report illustrates the use of ARTEMIS in various high-level applications, including the application interface using the cdev device support API above EPICS. Emphasis is placed on the design and implementation aspects of high-level applications which utilize the ARTEMIS server for information on beam dynamics.

Introduction

CEBAF is a 4 GEV electron accelerator facility which is in the final stages of commissioning. It consists of two 400 MeV superconducting linacs with a 5-pass beam recirculation system. The facility is capable of simultaneously serving three experimental halls with beams of differing energies.

The goal of delivering high quality beams to experimenters requires the availability of appropriate control, diagnostics, and monitoring functions to direct the complex operation of the accelerator. In addition, the efficient operation of a complex accelerator requires the automation of as many routine machine functions as possible. Automated computer algorithms are often assisted by a numerical representation of the system under control, or model. Many differing applications require the same modeling information, so centralizing the source of the various model data has many distinct advantages.

An early decision in the design phase of the control system for CEBAF was to base all high level functions involving machine setup and operation on accelerator models rather than resorting to a "look-and adjust" method of operation. This was deemed particularly crucial during the commissioning phase of CEBAF, when it was required to reconcile the machine behavior with its model.

I. ARTEMIS Model

The Accelerator Real TimE Modeling Information Server (ARTEMIS), currently under integration with operations codes, is a central server/client facility in the CEBAF accelerator control system, providing various model data (transfer matrices, twiss parameters, etc.) and supporting computations (e.g. quad strength calculation for matching) for all model-driven facilities. Centralizing the model calculations provides a uniform and consistent data collection mechanism for these and other applications, while eliminating the need for redundant calculations by different application software.

There are numerous optics calculations which are required for high-level controls. ARTEMIS currently provides the following optics calculations:

- First and second-order transfer functions for all optical elements at CEBAF.
- Lattice information (element position, type, etc.).
- Beam parameterization functions.
- Second-order ray tracing.

Calculations slated for inclusion in the near future include:

- Non-relativistic particle propagation.
- Spin polarization tracking.

One of the main objectives for ARTEMIS was to make available accurate accelerator lattice information on a timely scale, which implies the capability for the model to reflect changes in machine parameters at an acceptable rate. This led to the different update mechanisms for the model, which are periodic model update (time-based), event-triggered update (such as control system parameter modification) and user-initiated update.

The server/client facility of ARTEMIS is achieved by the use of cdev, a common object protocol currently under development at CEBAF which provides a standard interface to any data which can communicate with any source, including accelerator control systems, high-level applications, and now ARTEMIS. This integration of data representation is highly useful to the application developer and user, since this allows the origin of data to be abstract and of no consequence to the application and/or user. The facility of obtaining modeling operations, such as machine beta functions, by a standard mechanism is immediately evident.

* Supported by U.S. DOE Contract DE-AC05-84-ER40150

ARTEMIS is currently interfaced with several on-line facilities such as Tcl/Tk, allowing for the rapid prototyping of model-based algorithms and applications before they are made a permanent part of the control system, which will be illustrated below.

II. High-level Applications

There are several high-level applications which are being modified to use modeling information from ARTEMIS. The program currently under test is the CEBAF orbit and energy slow feedback lock, implemented in Tcl[1]. This code compensates for drifts in cavity RF phase and amplitude, which affect both beam energy and trajectory. Additionally, the code maintains a constant beam trajectory using a small number of correctors and monitors. The code requires the calculation of a response matrix propagating the vector (***x, x', y, y', dp/p***). The slow feedback control previously employed a Tcl-based distributed model database which serves as a prototype model.

The fast feedback control program [2], under development at CEBAF employs a modern control theory implementation of adaptive correction at a relatively fast update interval (presently 60 Hz). The optimal state estimator (determining ***x, x', y, y', dE/E***) will use response matrices from ARTEMIS for initial setup. Once the feedback process has begun, the adjusted response matrices can be compared with ARTEMIS for model verification and performance testing.

CEBAF also employs a Linac Energy Management program (LEM), which provides control over linac RF and optics to maintain a desired energy profile. ARTEMIS will provide LEM with an energy gain profile which reflects the current energy induced into the beam, taking into account cavity phase and gradient and relativistic phase-slip effects. Machine Twiss parameters from ARTEMIS will determine current accelerator optics, providing LEM with the information needed to efficiently maintain a desired energy profile and corresponding machine optics. LEM may also use "scratch-model" instantiations of ARTEMIS to predict what results will be achieved with any given change in energy profile and/or optics before committing changes to actual accelerator settings.

An on-line automated beam emittance and Twiss parameter measurement facility is currently in the design phase of development. CEBAF will utilize three different measurement techniques. The first method is a varying-optics method in which the beam is swept through a waist at a profile measurement device, using an upstream focusing element. ARTEMIS, in real-time modeling mode, will deliver transport information to the emittance calculation. The second method employs several profile measurement devices which will yield beam cross-sections at various points on the beam transport. As in the first method, the model server will provide transport information. The third method uses a thin "slit" aperture to determine beam position and slope. This information can be used to verify and/or update ARTEMIS model information. Once the beam parameters are measured, ARTEMIS will be utilized to perform a "back-propagation" to a selected origin in order to provide lattice matching.

```
blt_graph .new  -height 12c -width 22c
.new xaxis configure -title "Path Length (Meters)"
.new yaxis configure -title "Beta (Meters)"
pack .new -expand true -fill both
.new element create bx

set ElementNames [cdev model "get elements" "section arc1" "type {QUADRUPOLE SEXTUPOLES HKICK}"]

foreach Element $ElementNames {
   set bx [cdev $Element "get betax"]

   set i [cdev $Element "get scoord"]

   .new element append bx "$i $bx"
}
```

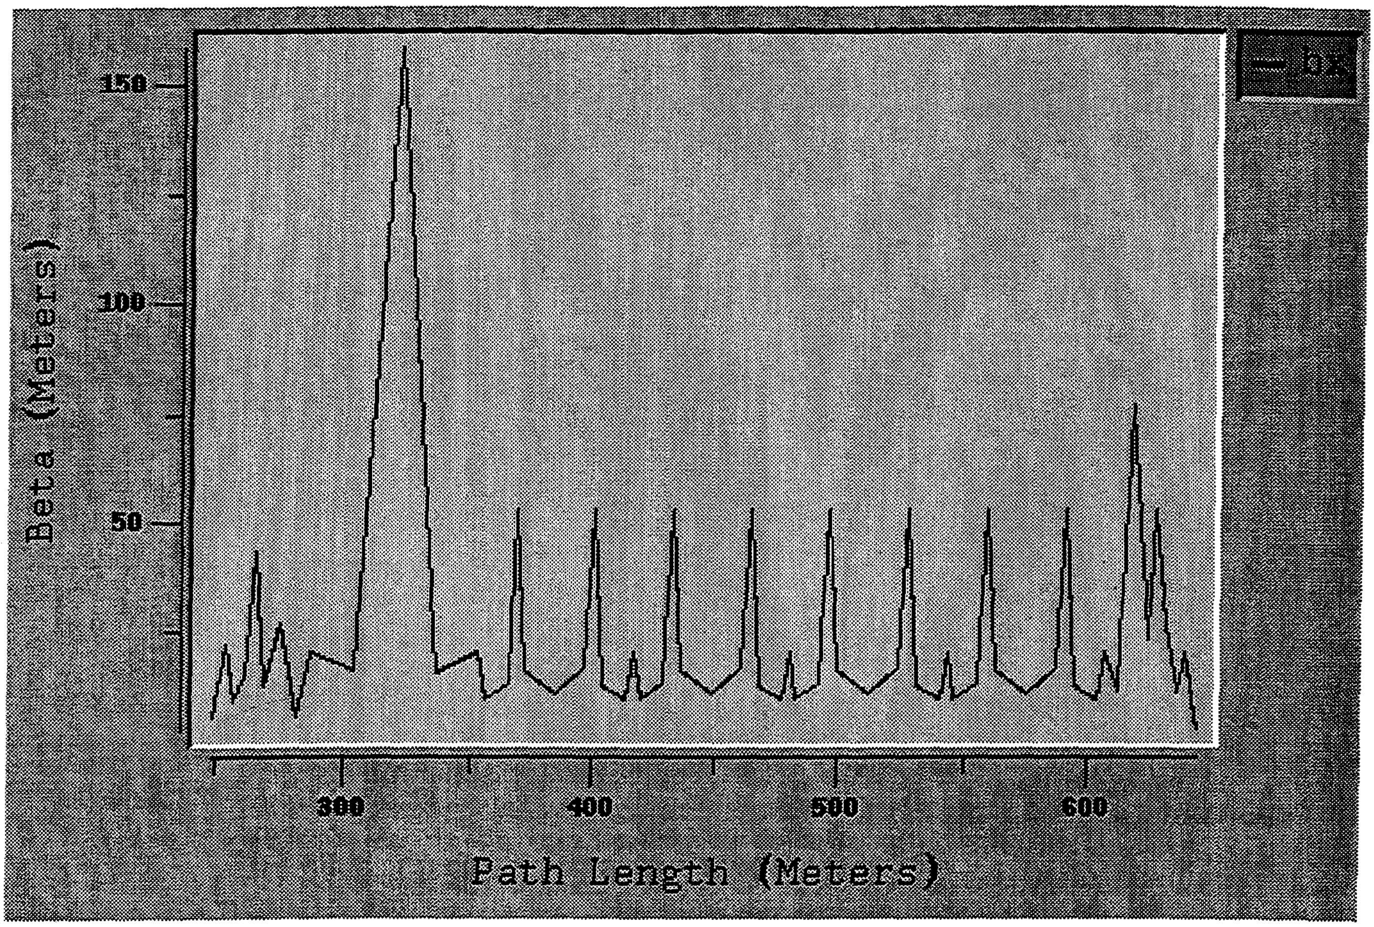

Fig. 1 Beta Function Plot Example

III. ARTEMIS Application Demonstration

ARTEMIS, with the integration of a scripting language such as Tcl, can be used to develop powerful applications with relative ease. For instance, graphical plots of machine beta functions as a function of path length is an operation which is performed often by accelerator physicists. To illustrate such an application, refer to Figure 1.

The first five lines of Tcl code are used to create the graph object with appropriate scaling and annotations. Line six is a query call to ARTEMIS, using the cdev data format, to retrieve the element names in the accelerator section **ARC1**, and of type **QUADRUPOLE, SEXTUPOLE,** and **HKICK** (inclusive OR), with the result returned in the variable **ElementNames.** This one function call initiates a connection with the ARTEMIS instantiation named **model** and performs the query. If the call fails, a callback procedure is executed, with the procedure defined by the application.

A loop is now defined on line 7 which loops over every element stored in ElementNames. The next line of executable code retrieves the horizontal (x) beta function for that element. This line is accessed using a reference to the element name, in which cdev performs the routing to access the model to retrieve the beta function value. The next line performs a similar query, this time obtaining the path length coordinate s. The final line appends the data to the plot variable **bx.** Execution time for the commands above (for 50 elements) operating on Hewlett-Packard 720 workstations (server on one workstation, client on the second, using ACE socket TCP/IP communications) takes approximately 3 seconds.

This application illustrates the two differing access modes to the server, which are direct model commands (element search) and access of data by requesting an attribute of a particular element (beta and path length commands). Nearly all of the commands to ARTEMIS support both access modes. The above example can easily be improved to enhance execution speed by eliminating the loop and requesting the beta function values as a single ARTEMIS command (similarly for the path length values). This will reduce network access overhead significantly, resulting in an execution time under one second.

The above example can be easily modified to perform other useful plots. For example, one can obtain real-time beam position information for all of the elements in **ElementNames** by using a single cdev call. These positions can then be divided by the square root of the beta function values and plotted, yielding a phase advance plot, a very useful optics tool. Using the callback mechanism, any changes in beam position or model parameters will automatically trigger an update to the graph.

IV. Summary

Many high-level applications at CEBAF are currently being converted to use the ARTEMIS accelerator model. The example above clearly demostrates the savings in application development and execution times which can be realized by utilizing ARTEMIS to obtain lattice and optics information. Future control applications at CEBAF will benefit.

V. References

[1] J. Van Zeijts, "Rapid Application Development Using the Tcl/Tk Language", (PAC 1995).

[2] M. Chowdhary, G.A.Krafft, H. Shoaee, W. Watson, "A Fast Feedback System for CEBAF", (ICALEPCS, 1995).

New Waveform Surveillance and Diagnostics for the LEP Injection Kickers

E. Carlier, V. Mertens and H. Verhagen
European Organization for Nuclear Research (CERN) - SL Division,
CH - 1211 Geneva 23, Switzerland

Abstract

The introduction of the Bunch Train Scheme in LEP requires a more precise and automatic supervision of the stability of the LEP injection kickers in timing and amplitude. Comprehensive and user-friendly diagnostic tools are required for in-depth investigation of equipment behaviour. A new system is currently being prepared using to a large extent commercial data acquisition hardware and hardware-independent software products.

1. EQUIPMENT

The LEP injection kickers [1] are fast pulsed magnets used to inject electron and positron bunches, arriving from pre-accelerators into the desired beam orbit in LEP. Three kicker magnets equally spaced along LEP, produce a fast orbit deformation of the stored bunches as shown in figure 1. At the moment when the kicker magnet deflects an already circulating bunch close to the field-free side of the septum, a newly injected bunch arrives at its field side and is bent by the septum magnet nearly parallel to the trajectory of the stored bunches. The duration of the fast orbit deformation must be sufficiently short to deflect only the stored bunch to which the new particles are to be added and must be exactly timed to avoid injection losses and residual beam oscillation.

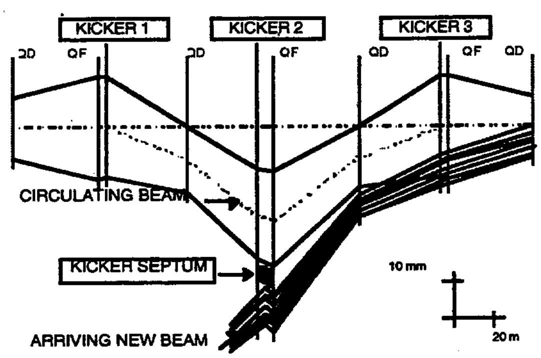

Fig 1. Layout of the fast orbit deflection for LEP injection, (QF, QD: focusing and defocusing quadrupole magnets).

The efficiency of the LEP injection kickers depends on two major parameters: the stability and reproducibility of the kick amplitude (better than 1 %) and the stability and precision of the kicker timing with respect to the injected and circulating beam (better than 25 ns). A typical kicker pulse, together with a signal of the injected beam current, is shown in figure 2.

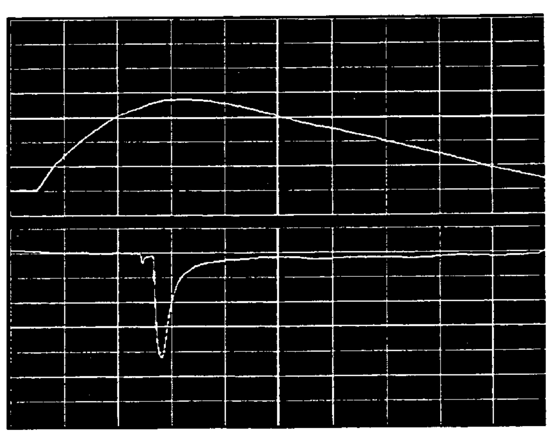

Fig 2. Typical kicker pulse (1) (300 A/div.) and injected beam current (2) (1 µs/div.).

2. SCOPE OF THE PROJECT

So far no permanent surveillance of the LEP injection kicker signals exists. The kick/beam timing synchronisation and the stability of the kick amplitude can only be controlled manually and only for following pulses, since no storage of previous pulses is foreseen. Faulty behaviour of the equipment can therefore only be confirmed if the symptoms persist. Monitoring of the evolution of the parameters in time to keep optimum performance is manpower-intensive.

Considerable improvements were nevertheless mandatory with the advent of the so-called 'bunch train scheme' [2] at LEP in 1995. In this scheme a much greater number of bunches circulate in the machine and a more precise injection scheme is needed. It has now become more important than before to precisely supervise and maintain the performance of the injection kickers and to trace back reasons of faults rapidly and with certainty.

It has therefore been decided to provide better surveillance and diagnostic facilities in the framework of an improvement project. These facilities comprise tools to select, acquire, log, retrieve and visualise signals from the beam pickup and the kicker system, to analyse them and to compare them with reference pulses. These tools should be either callable on request for diagnostic and adjustment, or run continuously in the background for equipment surveillance, with the possibility of informing the equipment specialists of any malfunctions. Finally, they could be used to readjust the timing and pulse heights automatically, if so desired.

3. GENERAL LAYOUT

In order to use commercial data acquisition hardware, to avoid as much as possible hardware dependence and to keep a high level of modularity, the application is being developed in the framework of LabVIEW [3].

The software part of the application is organized in a client/server scheme and divided into three layers: the data acquisition, the data analysis and the data presentation. The two first layers constituting the server part are running on a front-end computer located close to the acquisition hardware in order to reduce the data exchange over the network and profit from the event-driven capabilities of LabVIEW. The third layer, the client part of the application, is the graphical user interface which runs on HP workstations.

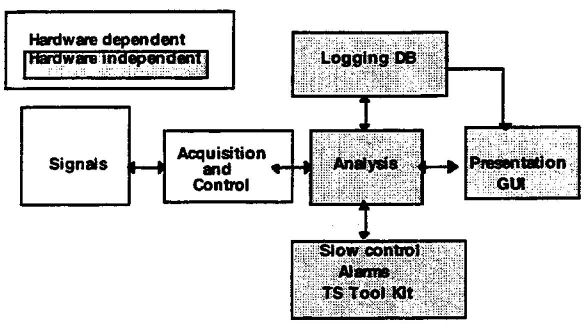

Fig 3. Software structure.

The acquisition layer includes the measurement devices and their integration. The hardware independence is achieved by using the LabVIEW virtual instruments driver library. The analysis layer comprises the tools used for signal surveillance and diagnostics. It also includes utilities through which the system interacts with other systems such as the general CERN alarm system, the equipment slow control, or databases. The presentation layer offers an uniform graphical user interface and look-and-feel for the application and can be run simultaneously from many computers without disturbing the acquisition and analysis processes.

The low level layers are seen by the top layer as virtual instruments. The communication between both levels is performed through the LabVIEW virtual instruments network based on BSD IPC using Internet stream socket [4].

4. ACQUISITION and DISPLAY

The kicker magnet, beam current and trigger signals are remotely selectable through multiplexer units and acquired through signal digitizers.

The single shot sampling rate is 500 MS/s for four channels with a resolution of 8 bits. The higher acquisition repetition rate is equal to 1.2s for a signal duration smaller than 2 ms.

The acquisition hardware will be connected to the existing accelerator control system (represented by a TCP/IP based Ethernet network) [5] either through a GPIB bus connected to a PC, plug-in DAQ boards integrated within a PC, or a VXI crate with a VXIpc card at the slot 0 position. In any case, DOS/MS Windows and LabVIEW will be used in the front end computer level.

The main advantage of the VXI solution is the existence of the VXI "plug&play" standard initiative which simplifies the setting up of VXI based systems. The same may hold true in the PC plug-in DAQ board solution when using the Windows 95 operating system. The GPIB solution offers the largest selection of instrumentation devices. All three ways of integrating the acquisition instruments are supported by LabVIEW.

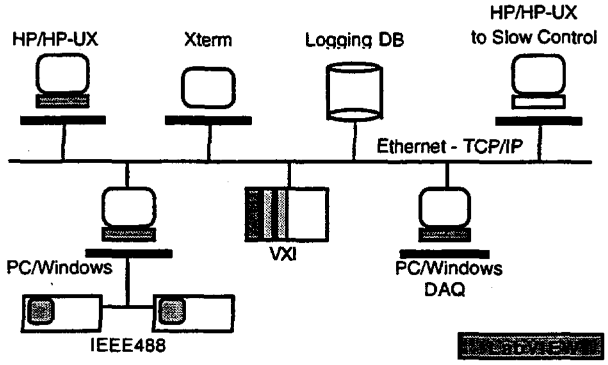

Fig 4. Overview of the high and low level acquisition system.

The waveforms are displayed graphically, on demand, on HP 9000/7xx series workstations, X terminals and PCs with X emulators, using the X Windows system.

A remote on-line Oracle database is used to store the default acquisition settings, the acquired and the reference waveforms. The connection to the kicker slow control is achieved through a fully configurable, event and data driven software package running on a HP workstation [6].

5. ANALYSIS

The selected waveforms (kickers + beam signals) are acquired upon reception of trigger prepulses when the system is pulsing and beam is available. The latest acquisitions are logged and kept in a local database organised as a FIFO. Upon request, the waveforms can be transferred and stored, together with the present timing and voltage settings, onto a remote Oracle database as reference waveforms.

Each newly acquired waveform is automatically compared with the corresponding reference. In order to detect differences in amplitude or in time, smart trigger facilities (envelope, level or window trigger) are used. If the new waveform is outside predetermined limits an internal trigger is generated and the last-measured waveforms are automatically stored with the reason for the fault in a remote Oracle database for later analysis.

Tools to select, extract, replay and display the acquisitions either from the local database or from the remote Oracle database are foreseen.

Upon detection of an internal trigger, a call to an external C function with predetermined parameters is made. This call connects the analysis layer to the existing equipment slow control software from where further actions can be initiated. Beside this interlock facility, different levels of alarms (warning, fault,..), depending on the severity of the deviation are provided. These alarms are sent to the general CERN alarm system through a call to the equipment slow control software [7].

For each new acquisition the difference in amplitude between the measured and reference signals is automatically checked at a predetermined number of points. The difference between the signals for the corresponding points is calculated and sent to the equipment slow control software for pulse to pulse regulation.

6. CONCLUSION

The project is actually in its design phase. Different technical solutions and configurations have been evaluated and it appears that the use of commercial hardware and software products for this kind of application tends to maximize flexibility, minimize obsolescence and reduce development time, effort and maintenance.

The use of well established, proven and long-lasting hardware and software leads to a minimum of specific development, even if the needs or the environment evolve, and this can reduce the cost over the lifetime of the system.

Due to manpower limitations, the realisation of this project will be outsourced to industry.

ACKNOWLEDGEMENTS

The authors would like to thank all persons who have increased their trust in the successful integration of commercial hardware and software products into the existing accelerator control system and those who have contributed with comments, suggestions, and information to the evolution of this project.

REFERENCES

[1] G.H. Schröder, J. Bonthond, U. Jansson, H. Kuhn, M. Mayer and G. Vossenberg, The Injection Kicker Systems of LEP, Proc. European Particle Accelerator Conference, Rome, Italy, June 7 - 11, 1988, Vol. 2, p. 1381, ISBN 9971-50-642-4, and CERN SPS/88-26 (ABT).

[2] B. Balhan, A. Burton, E. Carlier, J.-P. Deluen, J. H. Dieperink, N. Garrel, B. Goddard, R. Guinand, W. Kalbreier, M. Laffin, M. Lamont, V. Mertens, J. Poole and H. Verhagen, Modification of the LEP Electrostatic Separators Systems for Operation with Bunch Trains, Proc. 1995 Particle Accelerator Conference and International Conference on High-Energy Accelerators, Dallas, USA, May 1 - 5, 1995, CERN SL/95-45 (BT).

[3] National Instruments Corporation, Instrumentation Reference and Catalogue, 6504 Bridge Point Parkway, Austin, TX 78730-5039 USA, and World-Wide Web URL http://www.natinst.com/.

[4] National Instruments Corporation, LabVIEW Networking Reference Manual, part number 320587-01, 6504 Bridge Point Parkway, Austin, TX 78730-5039 USA.

[5] P. Anderssen, P. Charrue, R. Lauckner, P. Liénard, R. Rausch, M. Tyrrell and M. Vanden Eynden, Interfacing Industrial Equipment to CERN's Accelerators and Services Control System, CERN SL/95-42 (CO).

[6] V. Mertens, A. Aimar and E. Carlier, A Simple Generic Software Tool Kit for Distributed Controls Applications, Proc. International Conference on Accelerators and Large Experimental Physics Control Systems, Berlin, Germany, October 18 - 22, 1993, Nucl. Inst. Meth. A352(1994)427, CERN SL/93-49 (BT), and World-Wide Web URL http://dxbt00.cern.ch/ts/ts_papers/berlin93.hml.

[7] A. Aimar, E. Carlier, C. Cameron, A. Ferrari, B. Goddard, M. Laffin, V. Mertens and M. Tyrrell, Centralised Equipment Surveillance Using a Configurable Software Tool Kit as Front-End to the CERN Alarm System, CERN SL/Note 95-34 (BT).

An Introduction to Plant Monitoring Through the EPICS Control System*

J.A. Carwardine
Advanced Photon Source, Argonne National Laboratory
9700 South Cass Avenue, Argonne, Illinois 60439 USA

Abstract

The Experimental Physics and Industrial Control System (EPICS) environment [1] provides the framework for monitoring any equipment connected to it. Various tools offer engineers and scientists the opportunity to easily create high-level monitoring applications without having to rely on expert programmers to develop custom programs. This paper is aimed at the first-time or casual user, providing essential information for using several of the tools. Examples are taken from applications in regular use at the Advanced Photon Source (APS).

I. INTRODUCTION

Hardware interface to the EPICS control system is via local input-output controllers (IOCs), implemented at APS using VME-based 68,000-series microprocessors. Each IOC contains a database which references equipment connected to that IOC. Database records are referred to as *process variables* (PV), each one having a unique name. In general, a process variable record is associated with some hardware interface (I/O) connected to external equipment and contains the value associated with that I/O (e.g. the output value of a digital-analog converter, the input value from an analog-digital converter, the state of a binary input, etc.). Process variables can be analog, digital (two or more states), or entire waveforms, depending on the I/O configuration.

Access to process variable data (both input and output) is performed by "channel access," used by the EPICS top-level applications and is also available via C libraries to applications which need access to process variables. There are also utilities for command-line access (*caget* and *caput*). For example, the command *caget P:BM:DacAI* returns the value of the process variable *P:BM:DacAI*.

Each process variable record has a number of fields. In addition to the value field (the default), other fields provide information about the I/O; for example, the maximum and minimum values, the precision, error conditions, etc.

Operator interface to the control system is provided by top-level EPICS applications such as *medm* and the *alarm handler*. These have been supplemented by the SDDS toolkit [2] which can perform sophisticated data collection and post-processing functions.

II. FINDING PROCESS VARIABLE NAMES

The simplest way to find a process variable name is through a medm screen already containing the quantity of interest. By pointing to the quantity with the computer mouse and pushing the appropriate mouse button, the process variable name is displayed on the screen.

A recent addition to the APS control system is an application which allows wildcard searches for process variable names and also identifies the name of the IOC and its physical location in the facility. This application was written using the Tk/Tcl scripting language [3].

III. ALARM HANDLER PRIMER

The EPICS alarm handler [4] performs background monitoring of specified process variables and alerts the user to any change in state (e.g., if a power supply trips).

In order to use the alarm handler, a configuration file must be created. This lists the process variables to be monitored and how they are to be displayed to the user. The command *alh ex1.alhConfig &* starts the alarm handler with the configuration file *ex1.alhConfig* and produces a window such as the one shown in Figure 1.

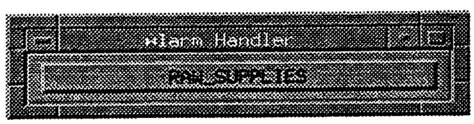

Figure 1: *Alarm Handler* Normal Screen

The color of the center part of the window indicates the overall condition of the process variables being monitored. If there are no alarms, the center is grey. If there are alarms present, the center may be white, yellow, or red (depending on the highest *severity* alarm present). If a new alarm occurs, the center of the window starts flashing (the color of the new alarm), and the computer starts beeping. Clicking the center part of the window with the mouse brings up a detailed screen (Figure 2).

*Work supported by U.S. Department of Energy, Office of Basic Energy Sciences, under Contract No. W-31-109-ENG-38.

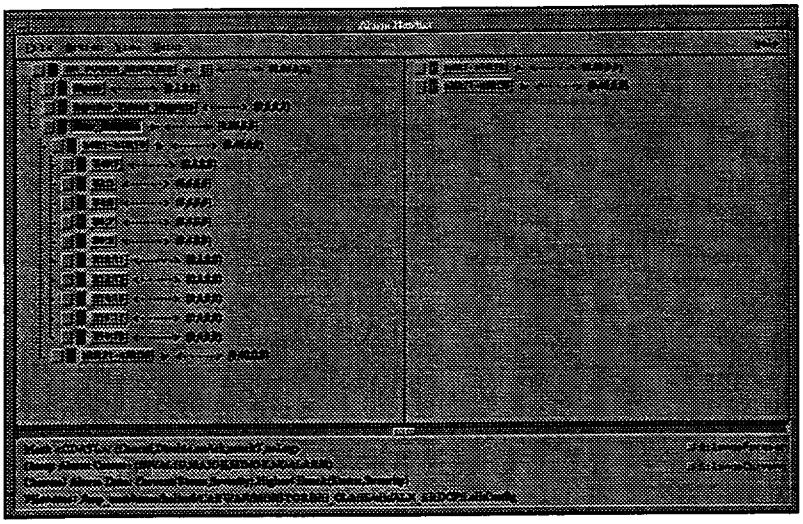

Figure 2: *Alarm Handler* Detailed Screen

This example shows the alarm handler detail screen for power supplies in the APS storage ring. Alarms are indicated by a small box of the appropriate color to the left of the process variable name. Alarms are acknowledged by clicking the box with the mouse (this does not reset the equipment). Unless the alarm has gone away, the alarm box remains, but the alarm handler stops beeping. When the alarm state goes away or is reset, the alarm box disappears from the detail screen.

A. Alarm Handler Features and Limitations

a) The alarm handler works with digital process variables. It therefore naturally lends itself to monitoring the state of a binary input (e.g., the status of a power supply). Analog parameters can be range-checked by setting up high and low limits in associated fields of the process variable, allowing the alarm handler, for example, to monitor that a power supply is not overheating.

b) Any one of three alarm severities (major, minor, invalid) can be associated with a process variable. 'Invalid' alarms are normally reserved for communications problems (e.g., with the IOC). Most other problems are configured as major alarms by default.

c) Process variables can be grouped for convenient display on the detail screen. This was shown in Figure 2, where power supplies were grouped by storage ring sector. The groups are shown on the left side of the screen in a tree-like structure. Highlighting a group shows the detail within that group on the right of the screen. Any number of levels can be created.

d) Buttons can be added to the alarm handler detail screen to allow the user to perform a specified action (e.g., pull up a detail control screen). It is also possible to configure the alarm handler to issue a command automatically if a given event occurs.

e) Other buttons can be added which display a specified text message when pushed. This can be used to offer guidance to the operator, for example.

B. Creating Alarm Handler Configuration Files

These can be created with any text editor (e.g., emacs). The following is the complete configuration file for monitoring eight process variables:

GROUP NULL RAW_SUPPLIES
GROUP RAW_SUPPLIES Sectors40&1
CHANNEL Sectors40&1 S40:1:R1:StatusCALC
CHANNEL Sectors40&1 S40:1:R2:StatusCALC
CHANNEL Sectors40&1 S40:1:R3:StatusCALC
CHANNEL Sectors40&1 S40:1:R4:StatusCALC
GROUP RAW_SUPPLIES Sectors2&3
CHANNEL Sectors2&3 S2:3:R1:StatusCALC
CHANNEL Sectors2&3 S2:3:R2:StatusCALC
CHANNEL Sectors2&3 S2:3:R3:StatusCALC
CHANNEL Sectors2&3 S2:3:R4:StatusCALC

The first command *GROUP* defines a new group called *RAW_SUPPLIES*. Contained within this group are two child groups called "*Sectors40&1*" and "*Sectors2&3*", defined in the two subsequent *GROUP* commands. The *CHANNEL* commands identify the process variables for the alarm handler to monitor. The formal usage of the above commands is:

GROUP <parent_name> <child_name>
CHANNEL <group_name> <PV_name>

Three additional commands are used to configure the auto-run commands and buttons which issue a specific UNIX command or display text messages. Their usages are as follows:

$ALARMCOMMAND <alarm> <unix command string>

Where *<alarm>* defines the event to trigger the issuing of the UNIX command string, e.g. UP_MAJOR, DOWN_ANY, would issue the command when a major alarm occurs or when any alarm goes away, respectively.

$COMMAND <unix command string>

This creates a button on the alarm handler detail screen next to the group or channel defined immediately prior to this statement. Pushing the button issues the UNIX command string.

$GUIDANCE <one or more lines of text> $END

This creates a button on the alarm handler detail screen next to the group or channel defined immediately prior to this statement. Pushing the button displays the text in a window.

IV. MEDM SCREENS

The EPICS application *medm* provides a graphical interface to the control system. Medm screens are the cornerstone of the operator interface to the control system and to the accelerator. Custom screens can easily be created graphically by running medm in edit mode (type *medm -edit &* from within an xterm window). The program has a similar look and feel to drawing packages for the PC or MAC and includes predefined graphical objects (such as slider bars, readbacks, etc.) for connecting to process variables. An example of a simple user-created medm screen is shown in Figure 3.

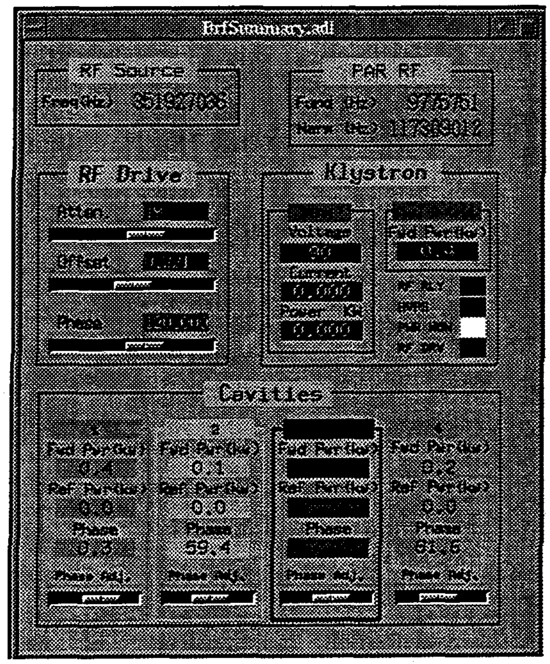

Figure 3: Example *medm* screen

V. SDDS PRIMER

The Self-Describing Data Set (SDDS) data format is widely used at APS [3]. Each data element in the file is given a unique ASCII name by the user which is subsequently used to access the data. An ASCII file header contains information about the data type and its location in the file. The user only need know the name of the data element to access it. Data can be integer values, floating point values or ASCII strings. Whilst the header is always ASCII, the data itself can be either in ASCII or binary format with no change in functionality.

There is a continuously expanding list of both general-purpose and custom programs which take advantage of the flexibility offered by the SDDS data format. For example, it is possible to perform complex mathematics, take statistics, do Fourier transforms, digital filtering, etc., and generate publication-quality plots using data contained in an SDDS file. There are also a number of EPICS-specific applications which collect data, automatically perform experiments, do closed-loop control, etc., all using the SDDS format. Tools also exist to convert data to/from other file formats, e.g., spreadsheets, oscilloscope data, proprietary mathematics packages, etc.

The SDDS tools are all UNIX command-line driven. This allows considerably more power to be incorporated into the programs than would otherwise be possible and allows the creation of custom data collection and analysis tools by combining the tools inside UNIX shell scripts. In some cases, visual interfaces have been created using Tk/Tcl to act as graphical front ends to the command-line programs.

A. The Minimum SDDS File

The following is a minimum SDDS file containing two parameters and two columns of data:

```
SDDS1
&parameter name="pi", type=double, &end
&parameter name="date", type=string, &end
&column name="value", type=double, &end
&column name="square", type=double, &end
&data mode=ascii, &end
3.1415928
3rd September 1995
7
0.0      0.0
1.0      1.0
2.0      4.0
3.0      9.0
4.0     16.0
5.0     25.0
6.0     36.0
```

The order of the data follows the order of the parameter and column definitions in the header. The number of rows in each column is indicated immediately before the start of the column data (in this case 7). The data type can be: *short, long, float, double, string.* A more comprehensive header might contain a text description of the file and many parameters and columns, each with a text description and units (e.g., Amps, mS, etc.).

B. Viewing SDDS Data

Clearly, small ASCII files like the example above could be viewed using either an editor or the UNIX command *more.* But in general, this really isn't a useful way of viewing the data. There are two parts to viewing a file: finding out what is in a file and looking at the data itself.

The command *sddsquery* allows the user to find out what is in the file. For example, if the above file were called *ex1.sdds,* then the command *sddsquery -column ex1.sdds* would show the names of all the columns in the file. Parameter names can be found by replacing *column* with *parameter.*

If more detailed information is required (say the data type), then the command *sddsquery ex1.sdds* would print out all of the information from the file header.

Data from the file can be accessed using the command *sddsprintout.* This allows printouts of data from a file for specified columns or parameters. For example, the command *sddsprintout ex1.sdds -column=value* will print all rows from the column *value* in the above file as follows:

```
Printout for SDDS file ex1.sdds
   value
---------------
 0.000000e+00
 1.000000e+00
 2.000000e+00
 3.000000e+00
 4.000000e+00
 5.000000e+00
 6.000000e+00
```

Wildcards can be used to select parameters or columns. For example, to view all of the column data, type *sddsprintout ex1.sdds -col='*'*.

C. Plotting SDDS Data

By far the most frequently used program in the SDDS toolkit is *sddsplot*. This is a powerful publication-quality plotting program for SDDS data files.

To plot the two columns of data in the above example file, type *sddsplot -col=value,square ex1.sdds*. Figure 4 shows the window which is created.

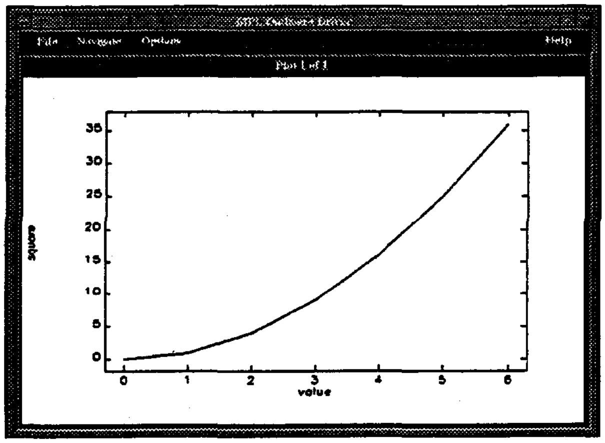

Figure 4: Window generated by *sddsplot*

The plot format can be modified in many different ways. For example, to plot the data points without joining the dots type *sddsplot -col=value,square -graphic=symbol ex2.sdds*.

More than one column can be plotted at the same time. Suppose a column *root* were added to the file *ex1.sdds*. Then both the *square* and *root* columns could be plotted together by using the command *sddsplot -col-value,square -col=value,root ex2a.sdds*. The same result could be obtained by combining the *square* and *root* column names into a single *-column* option with the command *sddsplot -col=value,'(square,root)' ex2.sdds* (note the use of single quotes and braces). Alternatively, the command *sddsplot -col=value,'*' ex2.sdds* would plot all the columns against *value*. There are many other ways in which the plots can be formatted. Examples can be found in [2] and [4]. Plots can also be formatted for inclusion in documents without resorting to an xwindows screen grab.

The following is a selection of the many options available in *sddsplot*:

-graphic allows the user to chose from connected lines, unconnected symbols, or unconnected dots.

-scale allows the user to specify the x- and y-scales.

-separate puts multiple plot requests on separate pages, e.g. *-col=value,'(square,root)' -sep* will produce two separate plots, one being *value* vs. *square*, the other *value* vs. *root*.

-samescale is used with *-separate* option, and forces all plots to have the same x- and y-scales.

-layout defines the layout of plot frames on each page, for example, plots could be put in two rows of two plots per page.

-topline adds a title to the plot.

-xlabel allows labelling of the x-axis with a supplied string, rather than using column names from the file. A similar *-ylabel* option exists for the y-axis.

-filenames adds the names of the data files to each plot.

-date timestamps the plot with the time and date.

D. A Note about SDDS Command-line Formats

All SDDS commands have a generic command-line format. The minimum command line would consist of the command plus a filename, although most commands require at least one option as well. All options consist of the "-" character immediately followed by a keyword (e.g. *-column, -ascii*, etc.).

Many keywords also require parameters, in which case the keyword is immediately followed by the "=" character plus the parameters separated by commas (e.g. *-column=value,square*). No spaces may appear within the expression.

All keywords can be abbreviated to the minimum unique character string (e.g., *-col* is recognized as being *-column*).

Any command-line argument not starting with the "-" character is assumed to be a filename. Many commands which modify an existing file can be given either one or two filenames. The first filename given is always assumed to be the input file and the second (if supplied) to be the name of the output file to be created. If only one is supplied then the input file is replaced.

In most cases, the order of the arguments does not matter. For example, *sddsquery -col ex2.sdds* and *sddsquery ex2.sdds -col* are equivalent. However, in some cases it does matter (e.g., the filename order described above).

E. Getting SDDS Help

All of the SDDS commands require some command-line arguments. Typing the SDDS command without any arguments gets a help message. Note that some can be very long (e.g. *sddsplot*). As an example, the command *sddsprintout* brings up the following message:

usage: sddsprintout [-pipe=[input][,output]]
[<SDDSinput>] [<outputfile>]
[-columns[=<name-list>[,format=<string>]][,endsline]]
[-parameters[=<name-list>[,format=<string>]][,endsline]]
[-array[=<name-list>[,format=<string>]]
[-fromPage=<number>] [-toPage=<number>]
[-formatDefaults=<SDDStype>=<format-string>[,...]]
[-width=<integer>]

There is also an online SDDS users' manual [5] which provides more detailed descriptions of the command usages.

VI. EPICS DATA COLLECTION

A. Using sddsmonitor

This program allows the value of any process variable to be logged on a periodic basis. For example, it may be required to monitor a power supply output current every 0.2 seconds for 10 minutes, or to log the storage ring vacuum every 10 seconds for 8 hours.

In order to use sddsmonitor, a configuration file must be created containing the names of the process variables to be monitored. Here is an example:

```
SDDS1
&description text="example list of names", &end
&column name="ControlName", type=string, &end
&column name="ReadbackName", type=string, &end
&data mode=ascii, &end
3
P:BM:CurrentAI     P:BM:Current
P:Q1:CurrentAI     PQ1:Current
P:Q2:CurrentAI     PQ2:Current
```

The first column, *ControlName,* is the process variable name; the second column, *ReadbackName*, is the name to be given to the data column in the output file. The *description* text string is optional.

Assuming the above file is called *ex1.monitor* and that a total of 20 samples are required, one every 2 seconds, the command would be *sddsmonitor ex1.monitor par.sdds -steps=20 -interval=2*. The new file *par.sdds* will contain the following columns:

Step
Time
TimeOfDay
DayOfMonth
P:BM:Current
PQ1:Current
PQ2:Current

The columns *TimeOfDay* and *DayOfMonth* are in units of hours and days, respectively.

Having collected the data, *sddsplot* could then be used to plot the data. For example, to plot the dipole current against time of day, type *sddsplot -col=TimeOfDay,P:BM:Current -topline="PAR Dipole Current vs. Time" par.sdds*. An example result is shown in Figure 5. Instead of specifying the number of steps, the total time for collecting data could be specified; e.g., in the above example, replacing the *-steps=20* by *-time=40* would produce the same result. The time can also be specified in hours or days, for example, *-time=2,d* collects data for 2 days, *-time=1,h* for 1 hour.

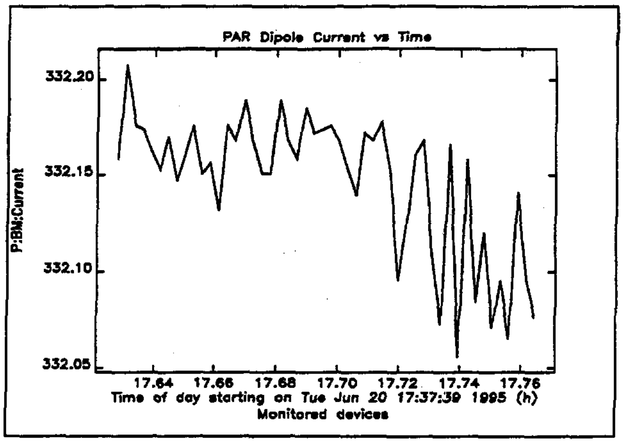

Figure 5: Example Data from sddsmonitor

Another feature of *sddsmonitor* is the *-trigger* option. This acts much like a digital oscilloscope trigger. Data is sampled at the specified rate but is not logged until the trigger event occurs. This is particularly useful for collecting data around a rare event. For example, it may be required to monitor the process variables in the above example, but only if the current exceeds 200A. This could be achieved by replacing *-steps=20* by *-trigger=B:BM:CurrentAI,level=200,slope=+,before=2,after=200.* A total of two samples would be logged before, and 200 samples logged after the current exceeded 200A.

B. Data Rate Limitations

Note that since the data is collected over the network, there are limitations to the rate at which data can be collected. Typically, a process variable cannot be read faster than about 10Hz.

C. Collecting Waveform Data

In some cases data is collected in the form of waveforms rather than individual samples and is processed by EPICS as a waveform process variable. A variant of *sddsmonitor*, called *sddswmonitor*, allows collection of data from these waveform process variables. The process variable names can be specified in a configuration file with columns supplied on the command line.

VII. SDDS DATA CONVERSION

Several programs are available to convert data between SDDS and other formats, either to analyze data from specialized equipment, or to make SDDS data available to other programs. Data can also be exchanged with several mathematical programs. Some of the available conversion programs are described below.

A. ASCII to Binary

In SDDS, ASCII and binary data files are completely interchangeable. The benefits of ASCII files are obvious, but as the files get larger, there is a size and access-time penalty, so binary files become preferable.

Conversion between the two data modes can be achieved using *sddsconvert*; e.g., to convert the file *ex1.sdds* to binary type *sddsconvert ex1.sdds -binary*. To convert back again use *-ascii* instead of *-binary.*

sddsconvert can also be used to rename or delete data.

B. SDDS to Spreadsheet

The program *sdds2spreadsheet* allows creation of an ASCII file for importing into spreadsheet programs. The resulting file can be delimited by any user-supplied character.

C. Oscilloscope to SDDS

Conversion routines have been written for a number of oscilloscopes and spectrum analyzers used at APS. For example, the program *tek2sdds* converts data from the Tektronics DSA602 oscilloscope to SDDS format. There are also programs to save and restore setups for various instruments using the SDDS data format.

VIII. UNIX SHELL SCRIPTS

The power of many of the above-described tools is significantly enhanced by the use of UNIX shell scripts to combine multiple tools, creating a new custom application. A shell script is a sequence of UNIX commands in a text file. The power of the shell script comes from the ability to quickly develop applications using existing programs and tools. They are extensively used by the APS operations staff for both machine control and analysis [6].

The simplest shell script might contain a single *sddsplot* command and be used to save typing on the command line, whereas a more complicated shell script may perform data collection and detailed analysis of power supply waveforms from the synchrotron accelerator. Any valid sequence of UNIX commands can be included in a shell script, and there are many programming-type features available such as program looping and the use of variables. The following example issues a series of caput commands to change the status of several process variables (in this case to open some vacuum valves).

```
#! /bin/csh
echo "This script opens all the gate valves"
echo "in the synchrotron"
echo "enter <yes> to continue"
if( "$<" != "yes" ) then
        echo "exiting"
        exit
endif
foreach valve (01 02 2A 03 04 05 06 07 07)
caput VM:BM:00"${valve}":openBO 1
end
```

Before the *caput* commands are issued, the script informs the user of its purpose and waits for input (indicated by the "$<" in the *if* statement). If the user does not respond with *yes*, then the program exits. Otherwise it issues nine *caput* commands, each time replacing *${valve}* with one of the strings 01, 02 etc.

The following example starts *sddsmonitor* with a predefined configuration file but with a user-supplied data file name:

```
#! /bin/csh
if( ${#argv} != 1 ) then
   echo "usage: startmon <outputfile>"
   echo "...Starts sddsmonitor for the par power supplies"
   echo "at 10 second intervals for 8 hours"
   exit
endif
set filename = $1
sddsmonitor PAR.mon ${filename} \
   -interval=10,s -time=8,h &
```

In this case, the script checks that the user has supplied a filename on the command line (the *if* statement). If not, a usage message is printed and the script exits. Otherwise the sddsmonitor program is started.

IX. SUMMARY

Using a combination of the EPICS control system, the SDDS toolkit, and UNIX shell scripts, users can easily develop custom monitoring applications with a minimum time investment and without having to develop a custom C program. Such applications are regularly developed by staff involved in day-to-day operations at APS. Often a few minutes spent writing a simple shell script can allow the control system to perform data collection and analysis, rather than having to perform the task manually. The examples given in this paper show that these are available to anyone with access to the control system and are not limited to expert software developers.

X. REFERENCES.

[1] Dalesio, L., Hill, J., Kraimer, M., Murray, D., Hunt, S., Claussen, M., Watson, C., Dalesio, J., "The Experimental Physics and Industrial Control System Architecture," Proceedings of ICALEPCS, Berlin, Germany, Oct. 18-22, 1993.

[2] Borland, M., Emery, L., "The Self-Describing Data Sets File Protocol and Toolkit," these proceedings.

[3] Ousterhout, J., Tcl and the Tk Toolkit, (Addison Wesley, 1994).

[4] Kramer, M., Cha, B-C K., Anderson, J., "The Alarm Handler Users' Guide," URL: http://www.aps.anl.gov/asd/controls/epics/manuals/alh/alh.html.

[5] Borland, M., "Users' Guide for SDDS Toolkit Version 1.4," URL: http://www.aps.anl.gov/asd/oag/manuals/SDDStoolkit/SDDStoolkit.html.

[6] Borland, M., Emery, L., Sereno, N., "Doing Accelerator Physics Using SDDS, UNIX, and EPICS," these proceedings.

Upgrade of the RF Control System in BEPC

Cheng Hui-jun, Zhai Tong-yan and Zhan Ming-chuan
Institute of High Energy Physics, Academic Sinica

1. INTRODUCTION

The Radio Frequency (RF) system is an important part of the Beijing Electron-Positron Collider (BEPC). It is used to accelerate the beams and maintain the energy of the electrons and positrons in the storage ring. Its correct operation is a precondition for the injection and storage of the beams.

Until February 1992 the BEPC RF system consisted of two cavities and the power for each cavity was supplied by four RF amplifiers. To improve the peformance of the collider and to meet the requirements for the luminosity upgrade, the RF system will be expanded by adding two more cavities and re-arranging the RF system. Instead of one cavity being powered by four amplifiers, in the new arrangement the amplifiers will be re-connected so that each cavity will be powered by two amplifiers. As a result, an upgrade of the RF control system becomes necessary, since its design is based on the two cavity mode and it cannot be extended to adapt to the new requirements.

The RF control system is required for switching the stations on and off, conditioning the RF fields in the cavities, logging of RF parameters and status, monitoring the state of local interlocks etc. The numbers of different control and monitoring points for an RF local station are shown in Table 1.

Type	DM status	ON/OFF control	Analog monitor	Analog control
Number	128	24	32	4

Table 1. Numbers of control points

After the upgrade, the RF control system can be divided into three levels: the Central Control level, the Local Control level and the Equipment level, as shown in Fig. 1.

2. CENTRAL CONTROL LEVEL

At the Central Control level, a VAX 4500 has been used as the host machine which controls all the local stations. A unit type 3922 serves to connect the VAX to a 2922 module in a CAMAC system crate. The CAMAC crate contains a Serial Branch Driver (SBD) with an Optical Data Link (ODL) module for each of the local stations.

3. LOCAL CONTROL LEVEL

The local control stations are situated close to the equipment to be controlled for convenience in connection, testing, conditioning and maintenence. At this level, the RF intelligent controller and the Programmable Logic Controller (PLC) form the kernel of the upgraded RF control system.

i. The connection between the Central Control level and the Local Control level.

Very strong RF and electromagnetic interference is caused by the high power equipment at BEPC. Therefore optical fiber data links are used for the connections between the Central Control level and the Local Control level. ODL modules, operating at 5 Mbit/s, have been developed at BEPC and are used for the CAMAC serial highway.

ii. RF local station CAMAC crate.

High performance CAMAC modules based on function are used. In this way, the number of types of modules is kept to a minimum and the design of the applications software for the local stations is simplified. Four types of general purpose modules are used. These are a 32-channel digital input module

(IDIM), a 32-channel digital output module (IDOM), a 32-channel smart analog input module (SAM) and a 16-channel high precision analog output module. A3 and A4 are input and output isolating amplifiers.

iii. RF intelligent controller and PLC

The RF intelligent controller is the major part of the upgrade and was designed and constructed at BEPC. There is one controller for each cavity. It is based on the 8031 single-chip microcomputer. It it deals with the signals of the control system and provides the following:

- on/off control of the low and high voltage supplies for the two RF amplifiers for each cavity
- analog signal input and display
- signal interlock and control of each cavity
- crucial signal aquisition and display

The schematic diagram of the RF controller is shown in Fig. 2

This, together with the PLC, also provides local manual control facilities. This facilitates testing, debugging and maintenance. The PLC manages the interlock logic so that, for instance, if the vacuum deteriorates in a cavity, it passes the information to the RF controller and the amplifier will be shut down. It takes into account several conditions and makes the appropiate decision. It also send signals to status and temperature displays.

4. EQUIPMENT LEVEL

At this level, the analogue and digital signals from the upper layers are received and acted upon and the required signals returned to the upper levels.

5. CONCLUSION.

The design of the local control station was started in February 1992 and the equipment was installed in October 1993. Since then, it has operated stably and reliably.

ACKNOWLEDGEMENT.
We wish to thank Bao Bo-yu for his help in the RF Control System design.

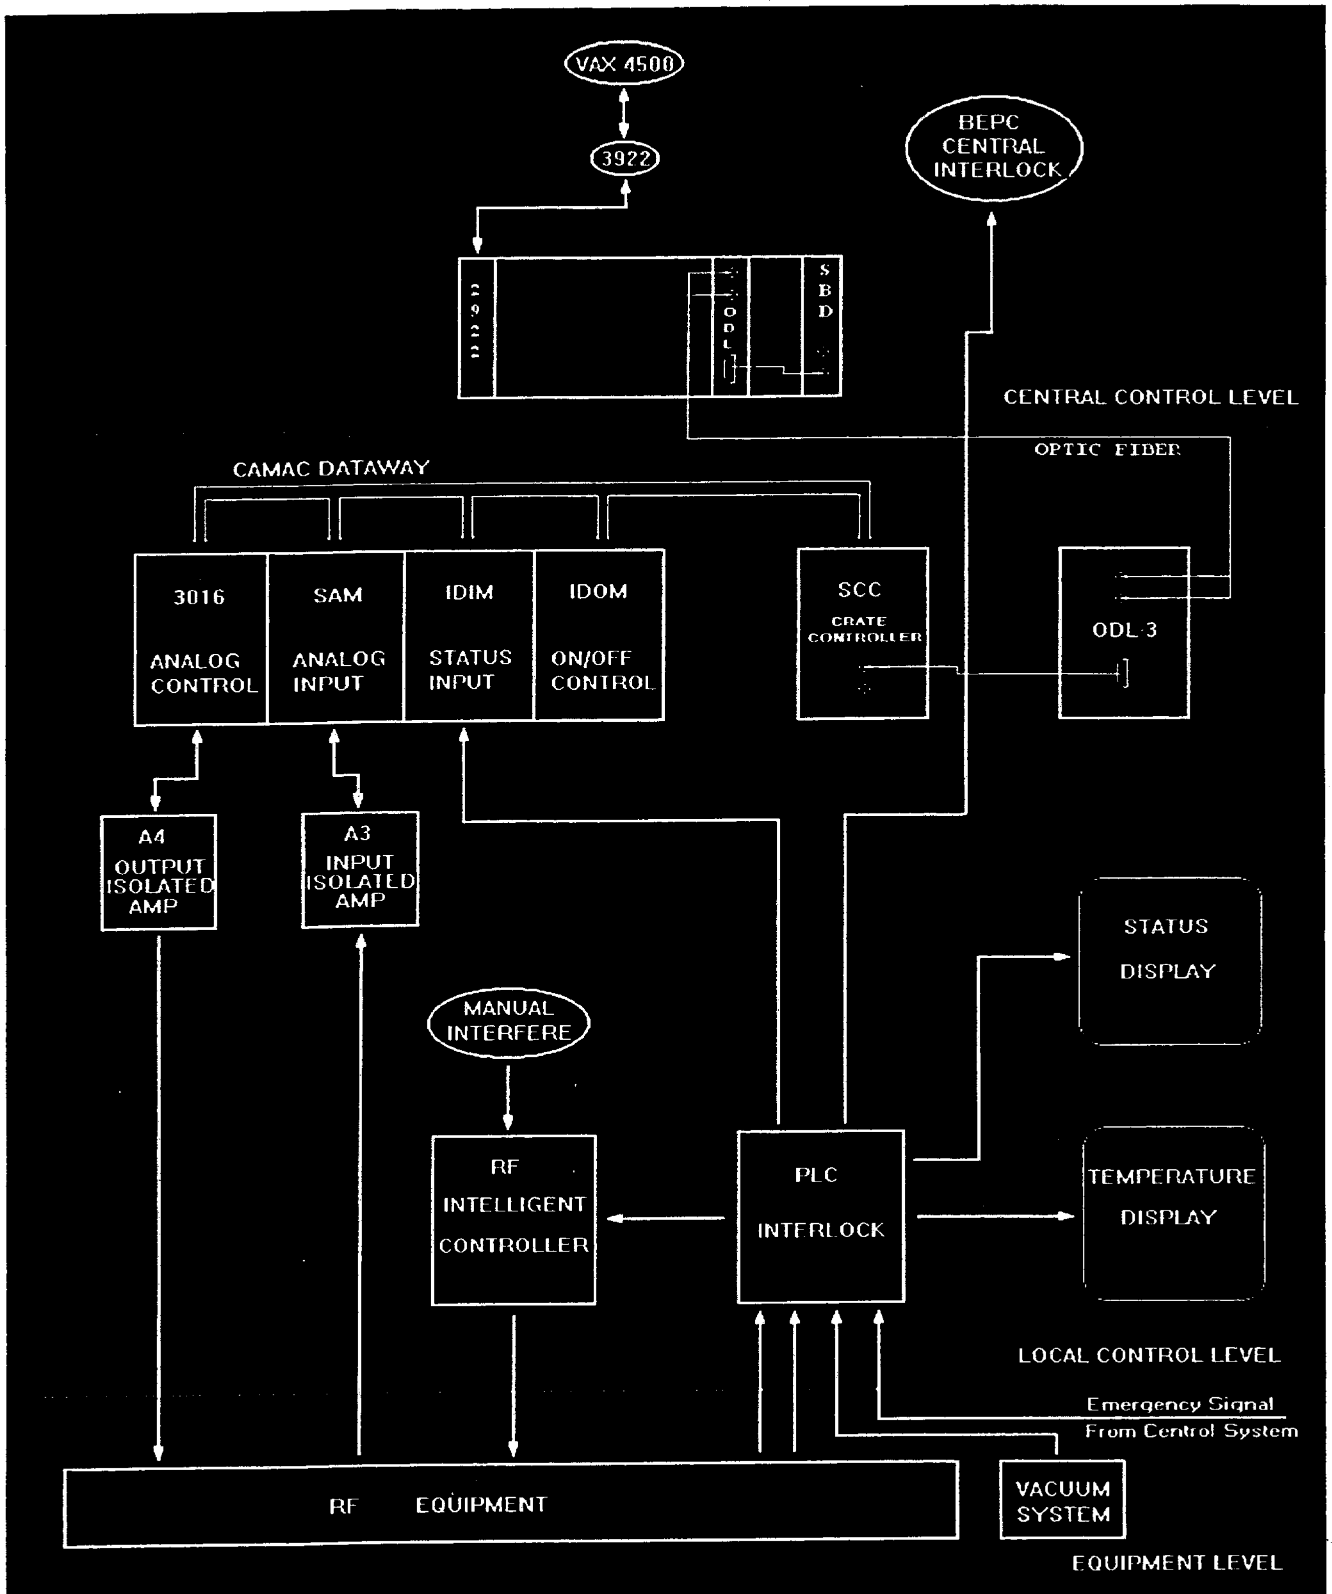

Figure 1. Block of RF Control System

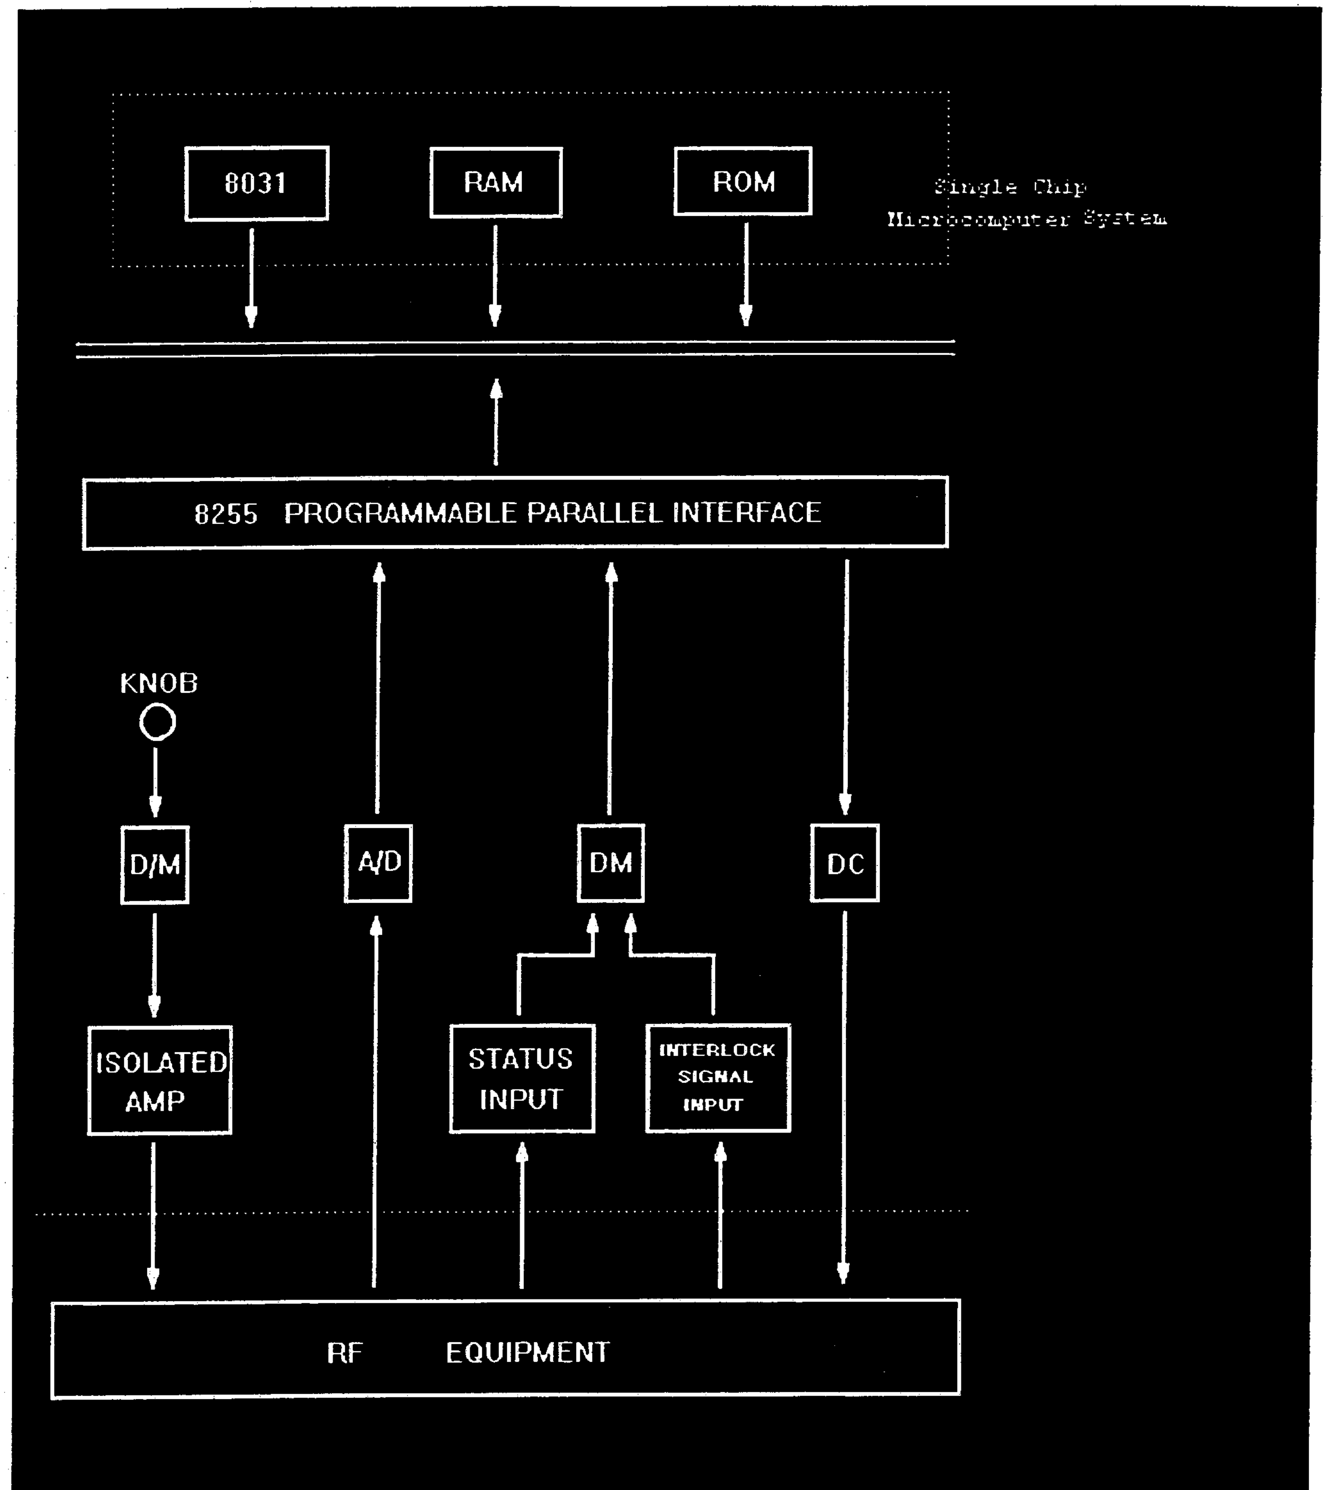

Figure 2. Block of RF Intelligent Controller

Computer System for Diagnostic of Electron Beam Using Transition Radiation

A.S.Chepurnov, A.S.Alimov, B.S.Ishkhanov, V.I.Shvedunov, P.L. Tkatchev

Institute of Nuclear Physics, Moscow State University.
119899, Moscow, Russia,
chas@rtm-cs.npi.msu.su, (095) 939-5659

Abstract.

A beam diagnostic system based on transition radiation analysis is a high-sensitive method for studying physical properties of the beam and adjusting of an accelerator. Optical transition radiation converted to a standard TV signal is analysed by the use of the unconventional computer based binary system of technical vision. It is possible to ensure beam diagnostics in real time due to the special hardware algorithms of image processing implemented in a video processor. This feature ensures cost effective solutions and simple software support. Data flows of digitised information are very compact and full of information. A diagnostic system is connected to the computer control system of a racetrack microtron injector - electron CW linac. Experimental results of the system application are described.

INTRODUCTION.

The optical transition radiation (OTR) effect [1-2] has gained general acceptance and can be used for many purposes involving beam diagnostics. Using the method based on trhe OTR effect one can receive all the necessary characteristics of a beam during the one cycle of OTR spot analysis:

- beam energy spread [3];
- beam positioning and beam profile [4];
- beam emittance [3, 5, 6];
- beam density and divergence [6].

The prime objective of our study was algorithms of image acquisition and processing, with the help of existing hardware. All parts of our image acquisition system are based on the LSI-11 compatible minicomputer. This computer type is widely used to-day for real time control of the injector of the Moscow CW racetrack microtron (RTM) [7, 8]. The system is connected with an IBM PC compatible computer (PC) via a standard interface. Graphical abilities of a PC are very suitable for man-machine interface support. The general scheme of the OTR analysis system is shown on figure 1.

The main feature of the computer control system is its control oriented architecture. Therefore, the presenting of a 3D image of the beam profile is an additional (not main) function of the image processing system. The main function of the system is "extraction" of digital indirect parameters corresponding to the beam properties. These parameters are used for closing feedback loops in the main control system to adjust output characteristics of the accelerator beam in real time.

Of course, in some operational modes of an accelerator it is very useful to observe an image of the beam in a way suitable for an operator. The operator could handle an overall image of the beam and "extract" simultaneously many parameters from the beam image, but the time of processing an image by an operator is not less than few seconds. But the question of adequate interpretation is very difficult and depends on the qualifications of the operator and the quality of the image restoring process.

The standard frequency of frames in normal TV signals is 50 Hz. Therefore, every 20 ms it is possible to get initial data for control algorithms for accelerator adjustments. In special applications when the beam has a pulsed nature with a frequency of more than 50 Hz, non-standard (higher) frequencies of frames could be used.

Commercially available frame grabbers are image oriented; they are oriented on the functional transformation, recording or storing of the whole frame. The algorithm of image processing for control application should be different. The main purpose of image processing for control is extracting a few digital parameters reflecting beam characteristics that could be used for digital closing of feedback loops.

The most simple example is the application of adjusting the electron optic system on the basis of the XY position of the beam OTR spot. XY coordinates of the beam spot center of gravity are only two numbers, and could be received as a result of a direct image processing system operation.

IMAGE PROCESSING HARDWARE.

OTR is converted to standard TV signals with the help of a TV camera. The image with the spot of OTR is stored and processed in an image acquisition board during the time of the next half-frame. A binary digitising system is used. A single bit sampling level could be set in a range of up to 256 levels of image brightness individually for every frame. The video processor unit supports in silicon the functions of square measurement, contouring, search on a raster and other special digital image processing functions. It is possible to realise beam diagnostics in real time due to features of the image acquisition board. Only necessary and useful data are transmitted to the control system, presented to operator and stored. The output data could be received in concrete beam parameters (size and position for example).

The hardware features ensure cost effective solutions and simple software support.

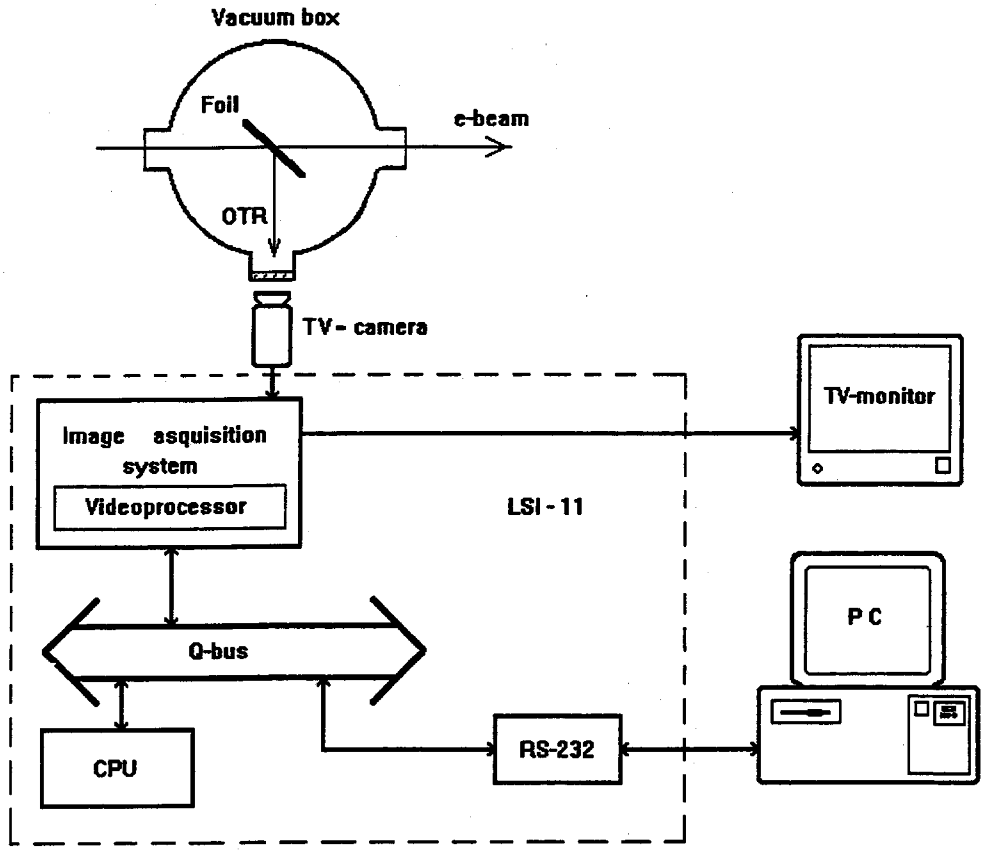

Figure 1. General scheme of OTR analysis system.

FIRST EXPERIMENTAL RESULTS.

The system has been used for studying an effect of the coupling slots on beam dynamics in the accelerating structure of RTM. This effect is explained by the transverse magnetic field excited on the axis of coupling cells and providing intercell coupling. Methods to compensate the coupling slot effect with external magnetic fields have been suggested [9]. For the exact measurements of the beam centre of gravity position and current distribution the method of OTR analysis has been used. At the axis of the beam 9 micron Al foil was placed at 45°.

Standard TV signals from a black and white CCD camera with a 14:1 magnification objective were analysed by a computer system. An example of a 20 serial frame 3D current distribution of the beam is presented in figure 2.

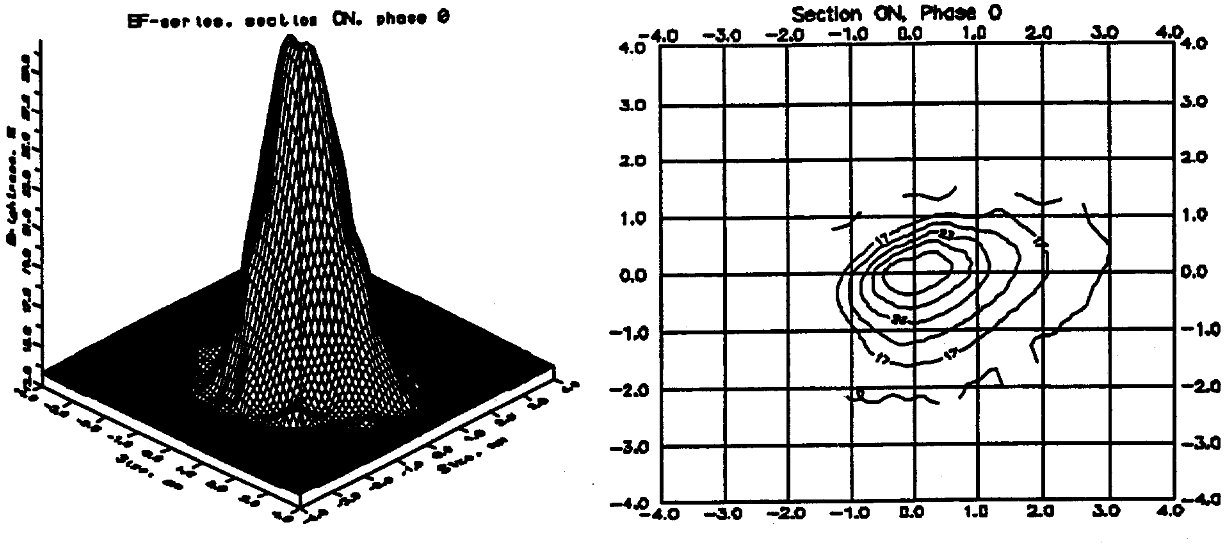

Figure 2. Restored 3D current distribution of the beam measured by means of OTR method.

IN THE FORESEEABLE FUTURE.

We are faced now with the problem of determining the accuracy of our measurement system carefully. The tasks of analog to digital channel linearization and increasing of the effective range and signal to noise ratio of the TV signal should be studied.

The image acquisition and processing algorithms will be implemented in a special hardware device based on modern technologies - digital signal processor and, soon FPGA. After the complete check-out a standard fieldbus interface (MIL-STD 1553B) will be added for fast connection with main computer control system of the accelerator, a system that is under construction now [10].

REFERENCES.

[1] I.Franc and V.Ginzburg, J.Phys. USSR 9, 353 (1945).

[2] L.Wartski, S.Roland, J.Lassale, M.Bolore and G.Filippi, J.Appl.Phys., vol. 46, No. 8, 3644 (1975).

[3] R.Chehad, M.Taurigna, G.Bienvenu, Beam Emittance Determination Using Optical Transition Radiation. Proceedings of the European Particle Accelerators Conference, 1977 pp.1139-1141.

[4] L.Wartsky et. al. Beam Diagnoctic Using Tramsitiom Radiation Produced by a 100 MeV Electron Beam. IEEE, Trans. Nucl. Sci, vol. 20, No. 5, 544 (1973).

[5] R.B.Florito, D.W.Rule, S.G.Iversen et al. Optical Transition Radiation Diagnostic for Low Emittance, High Energy, Charged Particle Beams, Proceedings of the European Particle Accelerators Conference, 1988. pp.1078-1080.

[6] S.G.Iversen, J.S.Ladish et. al. Charged Particle Beam Divergence Measurements Using Transition Radiation. IEEE CH2387-9/87/0000-0573.

[7] A.S.Alimov, V.I.Shvedunov, A.S.Chepurnov, O.V.Chubarov, et. al. Performance of the 6 MeV injector for the Moscow racetrack microtron. Nucl. Instr. and Meth. in Phys. Res., 1993, A326, pp.243-250.

[8] A.S.Chepurnov, I.V.Gribov et. al., Moscow University Race-Track Microtron Control System: Ideas and Developments //Proc. Int. Conf. on Accelerator and Large Experimental Physics Control Systems (Tsukuba, Japan, KEK, 11-15 Nov.1991.) Tsukuba, KEK, 1993, pp.140-142.

[9] V.I.Shvedunov, A.S.Alimov, A.S.Chepurnov et. al. Effect of Coupling Slots on Beam Dynamics in Accelerator Structure of Moscow CW RTM //Proceedings of the 1995 Particle Accelerator Conference, Dallas, USA, to be published.

[10] A.S.Chepurnov, A.A.Dorochin, K.A.Gudkov, V.E.Mnuskin, A.V.Shumakov. Family of smart devices on the base of DSP for Accelerator Control. (In these Proceedings).

New High Level Application Software for the Control of the SPS - LEP Beam Transfer Lines

F. Chevrier, G. Crockford, A. Hilaire, Y. Maumary, C. Niquille, V. Paris, P. Ribeiro, G. de Rijk, G. Smith
CERN SL Division, CH-1211 Geneva 23, Switzerland

New high level application software is being developed for the control of the SPS and LEP Transfer Lines. This paper briefly describes the model for the operation of these Transfer Lines, which is largely based on previous experience gained during the development and upgrades of the SPS and LEP control systems. The software system is then presented, followed by a description of the high level applications for the control room operators. Tools and methods used for the design and implementation of the system are mentioned.

1. INTRODUCTION

Transfer Lines

The Transfer Lines include all the equipment (dipole and quadrupole magnets, beam instrumentation, etc.) used to transfer the beams of particles into the Super Proton Synchrotron (SPS), and from the SPS onto targets or into the Large Electron Positron collider (LEP). They are divided into Transfer Zones, which join at intersections.

In order to run many experiments simultaneously, the SPS is used as a hadron accelerator and lepton injector for LEP in a time-sharing fashion. Because of the layout of the accelerator complex, some Transfer Zones are used for different types of particles (protons, electrons, positrons, Lead ions), in both directions, implying different settings for the steering or beam instrumentation. These settings are therefore functions of time, synchronised on the SPS Supercycle.

SPS Supercycle

The SPS accelerator is controlled in a so-called multicycling mode: at any one time the machine runs a repetitive sequence of beam injection(s), acceleration and extraction(s) of various particles. The whole repetitive sequence is called Supercycle and is composed of a number of elementary cycles, each of them divided into a beam segment and a preparation segment. From a functional point of view, each elementary cycle in a Supercycle can be considered as a separate accelerator.

The selection of a given segment to be worked on is done using a graphical representation of the Supercycle, the main bending magnets current.

2. TRANSFER LINES SOFTWARE

The analysis and design models for the operation of the Transfer Lines were largely based on previous work done for the SPS and LEP control systems [1,2]. See also [3].

The Transfer Lines software is organised around 3 ORACLE databases. The Offline database contains a full description of the machine in terms of hardware, as well as optics data. It serves as a source for the On-line Control database, which contains all the settings functions, hardware descriptions, hardware addresses and control parameters. The settings are organised into a hierarchy which allows control in terms of physics parameters, hardware magnitude or equipment settings, e.g. injection angle, dipole magnet strength or dipole current. A simplified data model of the On-line Control database is shown in Fig. 1. The system provides a reliable way of storing and retrieving settings, recording with a time-stamp all adjustments (trims) made, as well as providing a means of stepping back through these changes. All measurements are done by equipment Black Boxes, via the Measurement database. The latter contains measurement tables which are overwritten every time a new measurement is made and history tables filled from the measurement tables according to a trigger mechanism which can fire on different events (e.g. every measurement, on demand, etc.).

Taking into account the specific requirements for the operation of the Transfer Lines, the system was designed and structured into 5 subsystems, to be implemented as distinct applications for the control room. Fig. 2 represents a simplified model of the system. These applications exist and currently contain the minimum required functionality for operation. Another application is foreseen for the settings generation, which is currently being done manually (SQL

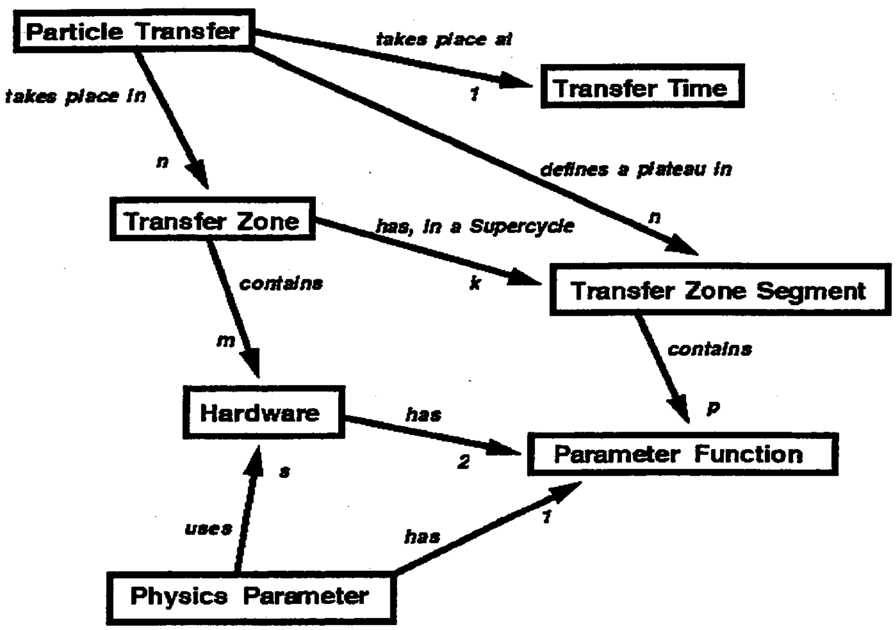

Fig. 1. Simplified Entity-Relationship Diagram for the On-line Control database.

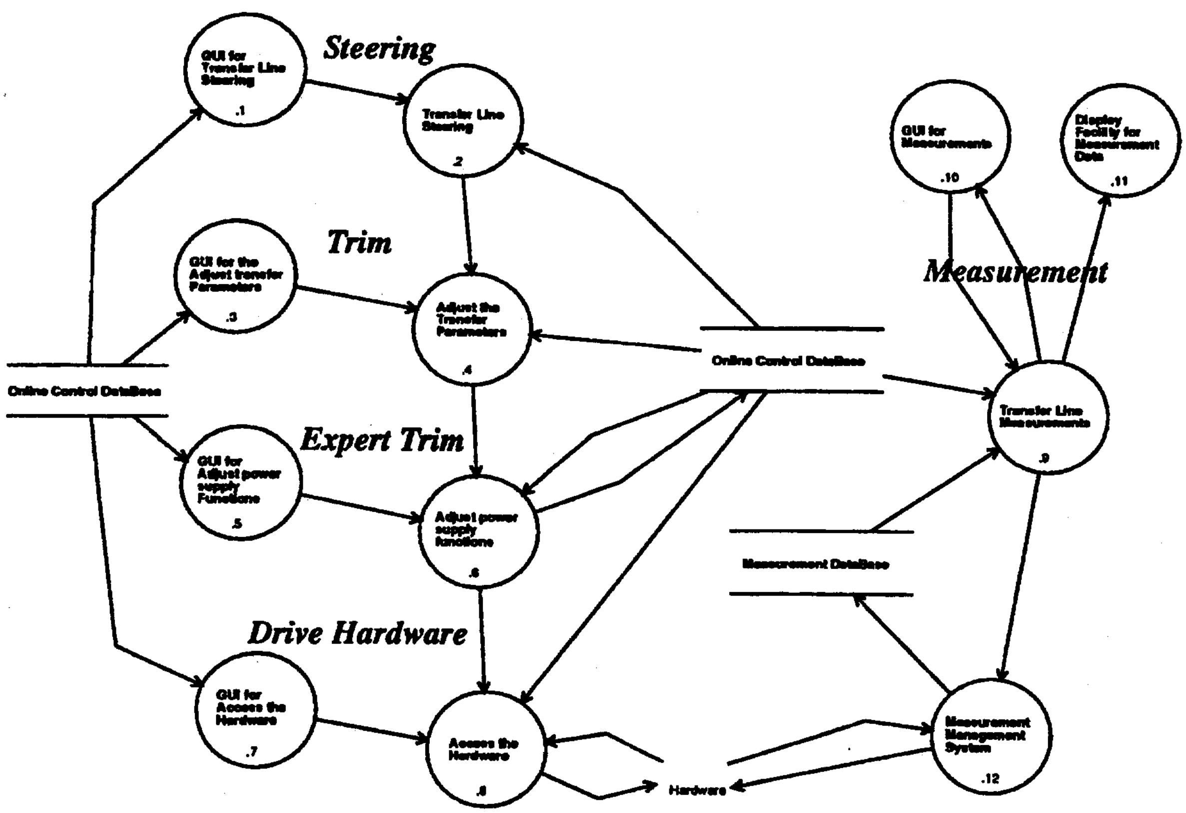

Fig. 2. Simplified model of the Transfer Lines software.

scripts).

3. THE APPLICATIONS

Throughout the design and implementation phases, emphasis was put on producing "data driven" software. To achieve this, the complete description of the Transfer Lines and their operation, in terms of which equipment is being used as well as optics data, is contained in the On-line Control database, allowing the control and management of a large amount of different equipment in a generic way. Furthermore, it gives the required flexibility which is especially important during machine development periods, when different optics and equipment configurations are often used.

To improve reliability and future maintenance, effort was also spent to produce modular software. The Graphical User Interfaces (GUIs) are written in X/Motif and were developed using the XUIMS [4] tool which helped to enforce modularity as well as uniformity in the "look and feel" through the use of templates. This uniformity is an important requirement for control room applications to ease the running of the accelerators and to minimize operator error.

Several other software packages, developed at CERN specifically for the operation environment, were also used [5].

The functionality of the subsystems is outlined briefly below.

Drive Hardware

This is a user-friendly facility to handle communications with the equipment in a generic way. It allows one to read settings from the database and load them into the hardware (low level equipment controllers), query hardware status, as well as perform other commands (e.g. initialisation, on, off, etc.). These operations can be performed on a single piece of equipment, a selected group, or even a whole Transfer Zone.

Measurements

The Measurements facility provides a standard interface to control the acquisition and display of all measurements of the beam itself or of the hardware. Measurements can be stored in the Measurement database for post run analysis and eventually on-line correlations. Graphical display is handled by the XPLORE package, which uses the commercial product XRT/graph (KL Group Inc.).

Trim

This application allows adjustments of the settings for a given Particle (beam) Transfer, or to revert to a previous state. The trimming can be done at 3 different levels: physics parameter (e.g. injection angle or position), hardware magnitude (e.g. kick strength in radians) or hardware (e.g. power converter current). It adjusts the value of the plateau corresponding to that Particle Transfer, in the parameter function. A full history of all trims is available for viewing, with the option of restoring any previous state, for a single piece of hardware or for all the equipment of a given type, for a given Particle Transfer.

Expert Trim

The parameter functions are described by a series of value-time coordinates, two consecutive points constituting a vector. The Expert Trim application allows a complete editing of these functions, at any level. It contains a graphical function editor including many specialized edit options (e.g. adding, removing or changing the value of a point, vertical shift of one or more vectors, etc.) that can be used in two modes: "click and drag" the mouse, or with keyboard input. As for the Trim facility, a full history of all trims is available, with restore options.

Steering

The Steering application allows correction of the beam trajectory of a beam transfer. It performs measurements of the trajectory, launches a correction algorithm to calculate adjustments and then sends them to the appropriate corrector magnets. The system allows for easy addition of new methods and several correction algorithms will be available in the future. The first method implemented is a manual adjustment of corrector magnet strength for designated detector-corrector pairs.

4. CONCLUSIONS AND FUTURE PLANS

At present, one Transfer Line (TT10, used for all beam transfers of protons, positrons and Lead ions from CPS to

SPS) is operated with the new system. This implementation is used to fine-tune the GUIs to the operators' requirements. The first conclusions that can be drawn from this experience is the need for simple interfaces for the frequent routine manipulations and the hiding of the rarely used, specialised functionality. This significantly contributes to faster and more reliable operation. Also, the "data driven" software design and implementation has proven to be a very worthwhile investment: many modifications in the control were possible without any modifications to the code.

The full implementation of all transfers, including injection and extraction into and out of the SPS, is planned for 1996. During the second half of that year, all special equipment (e.g. kicker magnets, electrostatic septa, stepping motor equipment) will be brought into the system as well. In parallel, higher level correction algorithms for the Steering will be implemented.

Selections of Zones or equipment are currently done through lists, which will be doubled by synoptic. Currently, only a global view of the Transfer Lines in the accelerator complex exists.

REFERENCES

[1] A. Ogle *et al.*, Experience with workstations for accelerator control at the CERN SPS, NIM A 293 (1990).

[2] R. Bailey *et al.*, Performance of the high level application software during LEP operation, NIM A 352 (1994).

[3] M. Lamont, What constitutes a good control system for operations, these Proceedings (ICALEPCS'95, Chicago, USA).

[4] P. Ninin, M. Vanden Eynden, XUIMS - The X-Window User Interface Management System, these Proceedings (ICALEPCS'95, Chicago, USA).

[5] M. Vanden Eynden, Display package XPLORE, CERN - SL/Note in preparation,
E. Hatziangeli *et al.*, Measurement Service MSRV, CERN - SL/Note 93-38 (CO),
M. Vanden Eynden, Software Synchronisation Mechanism SSM, CERN - SL/Note 93-71 (CO),
P. Charrue *et al.*, Equipment Access SL-EQUIP / RPC, CERN - SL/Note 93-86 (CO).

Control of Total Voltage in the Large Distributed RF System of LEP

E. Ciapala and L. Arnaudon

SL Division, CERN, 1211 Geneva 23, Switzerland

Abstract

The LEP RF system is made up of a large number of independent RF units situated around the ring near the interaction points. These have different available RF voltages depending on their type and they may be inactive or unable to provide full voltage for certain periods. The original RF voltage control system was based on local RF unit voltage function generators pre-loaded with individual tables for energy ramping. This was replaced this year by a more flexible global RF voltage control system. A central controller in the main control room has direct access to the units over the LEP TDM system via multiplexers and local serial links. It continuously checks the state of all the units and adjusts their voltages to maintain the desired total voltage under all conditions. This voltage is distributed among the individual units to reduce the adverse effects of RF voltage asymmetry around the machine as far as possible.

The central controller is a VME system with a 68040 CPU and real-time multitasking operating system. Event driven communication handlers allow fast reliable concurrent data communication with the remote units. The RF unit low level RF G64 equipment controllers use a VME 68030 CPU to achieve the necessary response time and reliability.

I. INTRODUCTION

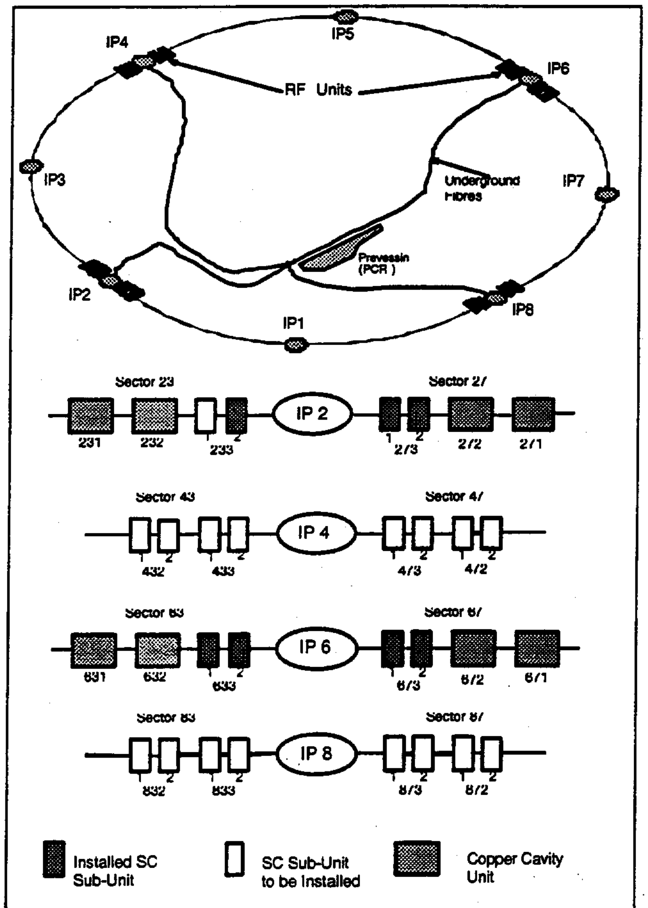

Figure 1. Present Distribution of the RF System in LEP

For the first phase of LEP and operation around 45.5 GeV, 128 room-temperature coupled-cavity assemblies were installed as eight separate RF units distributed around the Interaction Points (IPs) 2 and 6. Each unit consists of 16 cavities, two 1 MW klystron power sources operating at slightly different frequencies, high voltage power supply, auxiliary equipment, low level RF and controls. For the LEP2 upgrade to energies approaching 90 GeV 12 additional RF units based on super-conducting (SC) cavities, 192 in total, are at present being installed and commissioned. These new RF units make use of a similar infrastructure as far as possible. However both klystrons operate at the same frequency and each powers its own group of eight cavities. The RF on the two groups (sub-units) is therefore controlled independently. At the time of writing seven SC cavity sub-units are operational at IPs 2 and 6 together with the original copper cavity units. In the coming year a remaining sub-unit will become operational at IP 2 and four more SC cavity units will be commissioned at IPs 4 and 8. The total number of RF sources to be controlled will then be 32. The distribution of the RF system around the circumference of LEP is illustrated in Figure 1.

Operation of LEP requires that total RF voltage seen by the beam be maintained at a value which gives the required synchrotron tune Q_S at all times i.e. during injection, accumulation, energy ramping and physics running. At low energy, variation in Q_S produces harmful effects on the beam due to the crossing of synchro-betatron resonances. At high energy a value providing sufficient beam quantum lifetime must be maintained. At high energies it also becomes important to maintain symmetrical distribution of the RF voltage around the machine. The RF system is not evenly distributed around the machine and the various types of units and sub-units have different available minimum and maximum RF voltages. Furthermore particular RF units may be subject to field limitations or be in fault condition for certain periods. During running a unit may trip at any time due to any one of a large number of causes.

The method originally used for control of the total RF voltage was based simply on individually setting the voltage of each unit to get the required total. For energy ramping, when the total RF voltage must be increased to maintain Q_S, local RF voltage function generators were used. These were pre-loaded with calculated values prior to the ramp and triggered simultaneously by ramp events transmitted over the LEP general machine timing (GMT) system. This worked satisfactorily provided that all units continued to contribute the expected voltages during the ramp. It did not cope with unexpected changes in the state of individual RF units due to faults or interlock trips. In addition the calculation and loading of the individual ramp tables (300 values) for each new configuration was time consuming.

A single overall RF voltage control system is necessary. The problems in implementing this with classical analog techniques lie in the increased complexity of the local interface and the large distances involved. A computer based system with access to the RF units, monitoring their states and controlling individual voltages as required, can be made to provide sufficient performance with minimum RF hardware modification and can provide a high degree of flexibility. Initially a simple software system, running on a central workstation using the existing Ethernet and local GPIB connections to the low-level RF controllers of each RF unit was used to successfully maintain fixed voltage at injection or at top energy. However this was much too slow to follow the ramp with the necessary precision. Fixed bounds on the response of the system could not be estimated since both Ethernet and GPIB bus transaction times were dependent on other traffic and activities.

II. THE GLOBAL VOLTAGE CONTROL SYSTEM

An independent dedicated global RF voltage control (GVC) system based on a central controller with direct links to the equipment was first tested in 1993 [1]. This overcame the limitations described above and could maintain the RF voltage at its required value at all times, including the ramp. An improved version having better communications handling and increased functionality was made operational at the beginning of this year. This new system is also capable of adjusting the distribution of the RF voltage around the machine to avoid or at least minimize the effects of RF asymmetries on machine performance. Individual operating levels for the various units are now determined by a fixed strategy rather than by operator choice and this helps to provide better overall machine reproducibility.

General Layout

The general layout of the present system is shown in Figure 2. The central controller is situated in the Prevessin control room (PCR). TDM channels at 2 Mbit/s have been allocated for each of the points of LEP where RF is already installed (IPs 2 and 6) or will be installed (IPs 4 and 8). Serial line multiplexers at the end of each TDM link provide individual connection between the central controller and the RF units at each point. In the PCR connection is by sets of RS232 lines, one for each unit. In the underground klystron gallery where the distances are up to 500 metres the connection to the remote RF units is by RS422 differential transmission.

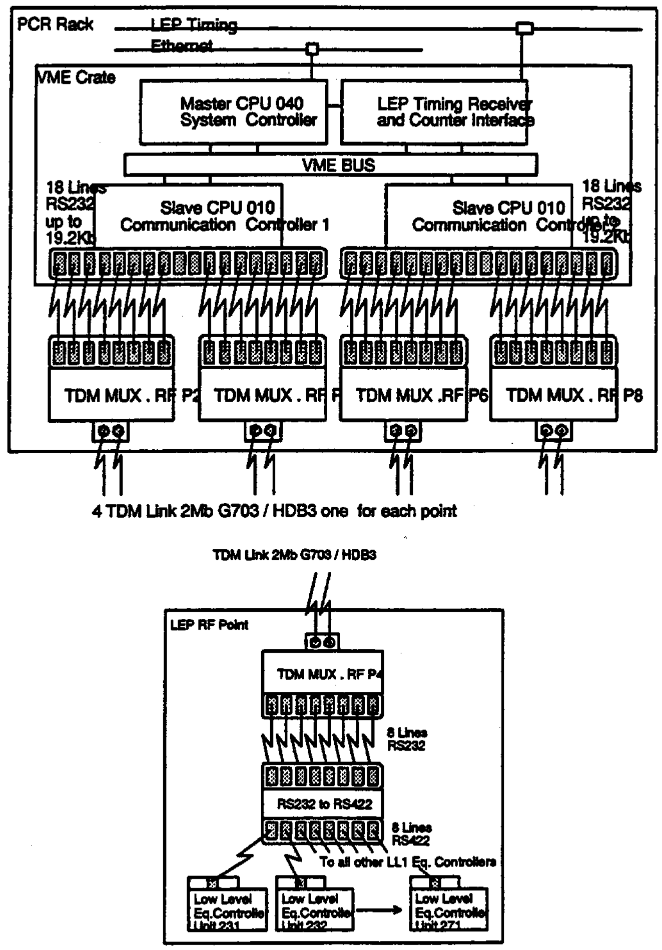

Figure 2. The Global Voltage Control System

GVC Controller

The central controller is a VME crate containing a 68040-based CPU module. Two slave 68010 CPU modules carry the multiple serial IO controllers for communication with the RF units. The OS-9 operating system is used. The global voltage control function is carried out by a C program which continuously monitors the state and RF level of each of the RF units or sub-units and applies corrections as required. Information on current unit states, operating levels (minimum and maximum) required voltage and symmetry method are stored in a local data module. Concurrent data transfer to the different units is implemented by having multiple communications handlers, one for each unit or sub-unit. These are in a sleep state until triggered simultaneously by operating system events sent by the main program when data is to be transferred. All data, previously received or waiting to be sent, is stored in the main data module. Other control and status acquisition programs read or set data in this module.

Data Transmission and Communication

The LEP time domain multiplexing (TDM) system is used for the dedicated connections needed to guarantee fixed maximum access time. This system uses fiber optic point-to-point links between the PCR and each of the IPs of LEP. It is used to carry machine timing systems and for other dedicated loop control applications. The lowest level provides 2 Mbits/s data channels with HDB3 protocol, the physical connection being to the CCITT G703 standard. Commercial equipment is used to further multiplex serial data channels within this bandwidth, 31 channels of 64 kbit/sec data being available. For each IP eight bi-directional serial channels presently operating at 9,600 Baud are used to connect the RF units and sub-units.

Low Level Equipment Interface

The serial RS422 line is connected to the low level Equipment Controller (EC) of the RF unit. This G64 hardware based EC contains the interfaces which allow the setting of RF voltage and phase, the reading of the RF sum of all cavity voltages and the state of the voltage control loop. The Low Level ECs in all units are fitted with a VME 68030 type processor module and VME to G64 converter [2]. Interrupt driven multi-tasking software is used. The global voltage control system can therefore set values and read information in the same way as for normal control via GPIB but independently of all other internal processes and with maximum priority.

System Operation

The main control program operates in a loop, continuously monitoring the state and voltage levels of the active RF units. This is shown in detail in Figure 3. If a new total voltage reference is required the voltages on the individual units are incremented or decremented accordingly. If a change in state of an active unit is found the voltages are redistributed in order to maintain the required total. The system is not a closed loop in the normal sense, since the actual voltage read by the cavity voltage detectors on the RF cavities is not used, but instead the current setting is used, together with a status bit indicating correct operation of the unit. This was done to avoid problems of

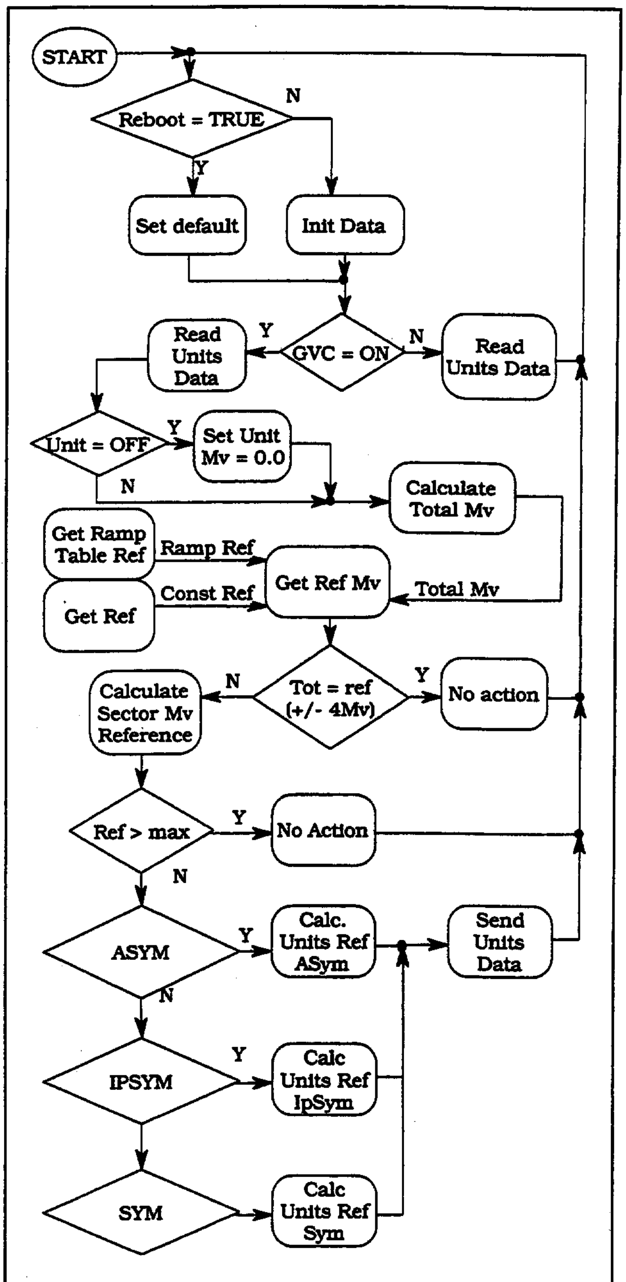

Figure 3. GVC Program Flow

instability resulting from noise and field oscillations on the RF units. The two modes of operation are defined in terms of the source of the voltage reference :

1) CONSTANT mode

This mode is applicable to injection and coast. The reference is a value of total RF voltage in MV. The operator selects the Q_S value required. Sloppysoft application software calculates the corresponding voltage reference for the GVC system taking the relevant machine parameters and settings into account.

2) RAMP mode

At the beginning of the fill the total RF voltages required at each 0.125 GeV step in energy during the ramp are loaded into a table in the central controller. The reference voltage during the ramp is derived from this table using a timer started and stopped by the ramp events sent over the GMT. At the end of the ramp CONSTANT mode is selected and Q_S trims are made if required.

Operations Interface

The system is set up and controlled for normal operation by the standard LEP operations software package, 'Sloppysoft'. Information about which units are to read, which units may be acted on, their available levels and other relevant information, is stored in an ORACLE database. This data, known as the RF Current Data Set (CDS), is used by all applications programs which deal with the RF system. The CDS is transferred to the GVC system when the running configuration is changed and during the setting up procedure prior to each LEP fill.

An X-window-based application, using a commercial graphics package, has been written to monitor the operation of the system and to display the current RF voltage and its distribution around the machine. An example of the data display is shown in Figure 4. This shows symmetrical distribution of RF voltage at each side of IPs 2 and 6.

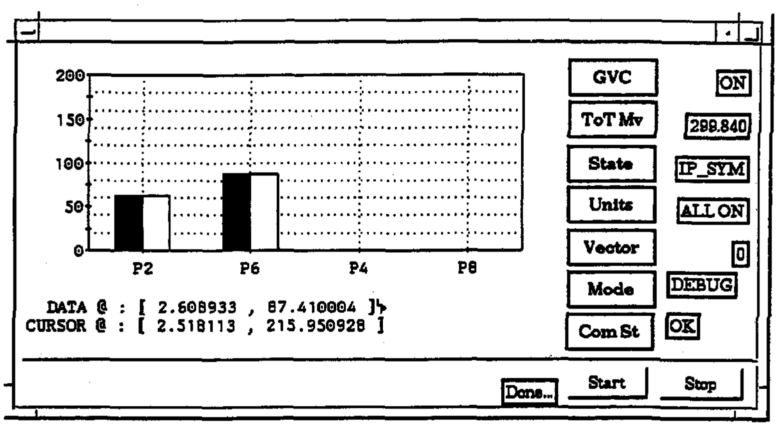

Figure 4. Display of Voltage in RF Sectors.

III. CONTROL OF TOTAL RF VOLTAGE AND RF ASYMMETRY

The total voltage must be maintained at a value which avoids synchrotron resonances (important at injection) and provides sufficient quantum lifetime at top energy. Quantum lifetime decreases rapidly once the RF voltage decreases below a certain value as shown in Figure 5 and the RF system must be run with sufficient reserve (80 MV at least) in order that the GVC system can restore the effect of a unit trip before other trips occur.

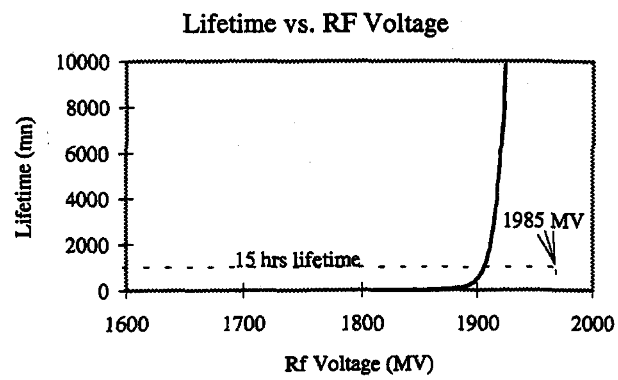

Figure 5. Quantum Lifetime vs. RF Voltage at 88 GeV.

The effects of IP RF asymmetry on machine operation and performance are mainly related to differences produced between electron and positron energy variations around the machine. Synchrotron radiation losses will be enormous in LEP2, around 1730 MV per turn at 88 GeV and this represents a significant fraction of the total beam energy. There is a gradual, almost linear, decrease in energy as the particles traverse the arcs. Energy is restored by the RF units around the RF equipped IPs, by an amount $V_{RF} \sin Ø_S$ where V_{RF} is the RF voltage amplitude and $Ø_S$ is the synchronous phase angle. This effect is known as the energy sawtooth. The shape is directly dependent on energy and on the RF distribution around the machine.

Examples of different energy sawtooth variation for different RF distributions are shown in Figure 6 for the simple case of RF installed at IPs 2 and 6 only. With completely symmetric RF distribution electrons and positrons have equal and opposite energy (and position) variations from the center. With IP 6 providing more RF voltage the amplitudes are increased, there are energy differences between electrons and positrons at IPs 4 and 8 but the center-of-mass collision energy is the same. Similarly if RF voltage is not equal on both sides of a given IP the same effect occurs. The IPs are equipped with low beta insertions which have strongly focusing super conducting quadrupoles. The insertion can only be correctly matched for one particular energy and errors produce modulation of the beta functions, phase advance differences in the cells for either or both particle types and a vertical tune difference (Q split). These effects are difficult to handle in operations and result in lost luminosity and background problems. The RF voltages around the machine must therefore be set to minimize these effects as far as possible under all conditions.

The GVC system can be operated with various conditions on symmetrical RF distribution. These can be specified by the operator and the system will maintain these as long as the available voltage at each point or individual unit permits. After limiting values for symmetry have been reached the RF voltage at each sector or on each unit is simply increased by an amount proportional to that which remains. The following conditions are allowed for :

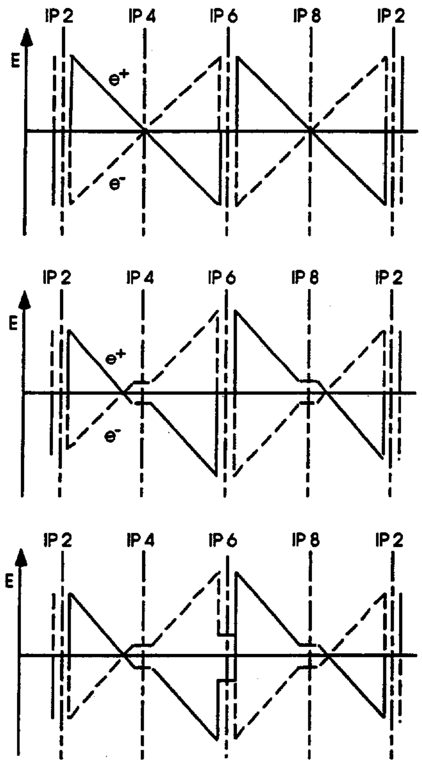

Figure 6. Energy Sawtoothing and RF Asymmetry.

- Asymmetrical - All units maintained at the same fraction of their individual maximum voltage. This allows ramping to maximum voltage with all units reaching maximum at the same time, without symmetry considerations

- IP symmetry - Equal voltages from RF on either side of the IP, all units at each side having the same fraction of their maximum.

- Symmetrical - Equal voltages at opposite interaction points, This can be with or without IP symmetry.

During 1995 the RF system was run with IP symmetry, to minimize Q split between electrons and positrons in the interaction points equipped with RF.

IV. SYSTEM PERFORMANCE

The effect of switching an active unit off is shown in Figure 7. The detector sum of all units was logged at 2 second intervals using a commercial instrument control package independently of the GVC system. Unit 231 tripped just before the start of the ramp. The system increased unit 232 to maximum, decreased sector 27 to maintain IP symmetry and increased all units at IP 6 to maintain the reference total voltage. This was done within one logging interval. The time of response to a change in RF reference is dependent mainly on the time taken to change the RF units to the new values. The normal ramp rate for the RF unit is 10% of maximum per second, but this can be increased to 40 % per second. The complete cycle of data acquisition, RF unit voltage recalculation and initiation of the voltage changes required in the RF units and sub-units is of the order of 200 ms. If an RF unit trips and Q_S is not too close to a harmful resonance or near the limit for very low quantum lifetime the voltage will be restored before critical beam loss. The overall response is also sufficiently fast to allow the voltage ramp function to be followed with an undetectable error in Q_S.

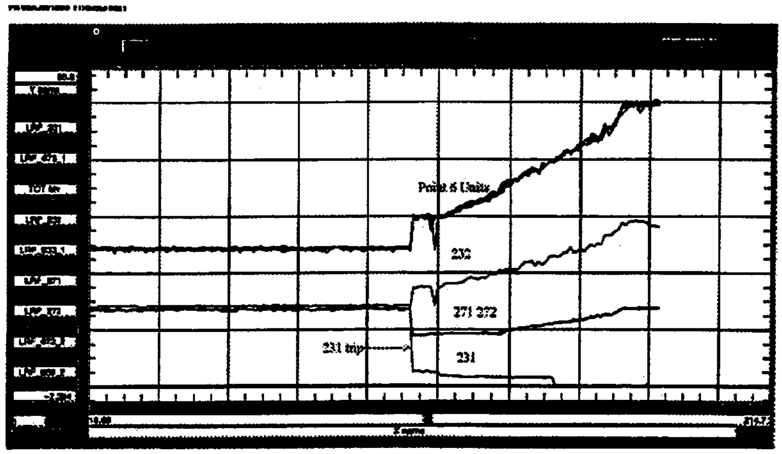

Figure 7. Response to Unit Trip at Start of Ramp

V. Improvements for Higher Energies

For running with SC cavities it may be necessary to de-tune certain RF cavities and perhaps adjust RF unit phases during acceleration. This can be done by the GVC system which can issue direct commands directly to the low level equipment or over the general control network to other equipment where time is less critical. To improve performance and to simplify operation the SC cavity units will be equipped with a fast RF feedback loop operating on klystron drive level instead of modulating anode voltage. This will allow a much faster ramp rate and the limitation in overall response will then be that of the GVC system. The speed of the system will therefore be increased during the shutdown by increasing the baud rate from 9600 to 19200. Individual commands sent to the RF unit will also be grouped to cut down the number of transactions.

A further general improvement would be the replacement of the RS232/TDM based communication system with direct links. The available options are being evaluated.

VI. Conclusions

A global RF voltage system has completely replaced the original local function generator based system. It was made operational for the 1995 startup and used routinely throughout the running period. It has successfully prevented beam losses due to RF unit trips. It has maintained symmetric RF distributions and helped to provide improved and more reproducible beam conditions. The addition of new RF units for LEP2 is straightforward and the system is able to handle all future requirements which can be envisaged at present. Software improvements will continue to be implemented and high speed data links evaluated for the replacement of the multiplexed TDM system.

VII. Acknowledgments

The TDM channels were provided by the SL-CO group and we are grateful to the specialists concerned for their assistance. The addition of GVC control in LEP operations software was made by Paul Collier.

VIII. References

[1] E. Ciapala, A. Butterworth, E. Peschardt, "Global Voltage Control for the LEP RF System," Proceedings of the 1993 Particle Accelerator Conference, Vol. 3, p1903 (1993)

[2] E. Ciapala, L. Arnaudon, "Use of VME CPU Modules for the Performance Upgrade of G64 Equipment Controllers for the LEP RF System," Proceedings of the Fourth European Particle Accelerator Conference, Vol. 3, p1833 (1994)

Large Scale Experience with Industrial Stepping Motor Controllers and Resolver Read-out Systems at SPS and LEP

R.J. Colchester, J.J. Gras, R. Jung, J. Koopman, J.M. Vouillot
European Organization for Nuclear Research, CERN, CH 1211 Geneva 23
J. Feres, B. Lopez
Midi Ingenierie, Labège Innopole, BP 131, F 31 676 Labège, France

Abstract
Over three hundred stepping motors and two hundred resolvers are used for Beam Instrumentation purposes in the SPS and LEP rings and transfer channels. Most of these instruments are beam intercepting and therefore require very reliable drive electronics. The coincidence of the upgrading of the SPS controls in 1993 and of the energy upgrade of LEP (LEP 2 project) was the ideal opportunity to standardise the motor control systems for both accelerators and to acquire an industrial system which would fulfil the requirements for the two machines in the most economical way. The interface to the control system had to be compatible with different environments, currently PCs in the SPS and VME crates in LEP and probably other systems in the future. A major problem encountered with industrial stepping motor drivers is the high level of EMI noise generated by the chopper type power converters, this had to be reduced drastically for this application. The project was handled as a "farming-out" project with CERN producing a detailed functional specification and the hardware development being entirely the responsibility of industry. A company was selected, and after extensive discussions, an order was placed to cover all SPS and LEP 2 needs. The first set of 210 motors and 80 resolvers has been equipped with the new control systems and it has been running for more than one and a half years without failures. Details of the hardware, software and EMI measurements, as well as present experience with the system are presented.

1. INTRODUCTION

The SPS accelerator and the LEP storage ring run for approximately 4000 hours per year, 24 hours a day, for uninterrupted periods of several months. Over 340 stepping motors are used in the SPS and LEP for Beam Instrumentation applications. These motors drive a variety of instruments, ranging from heavy Tungsten blocks used for beam collimation purposes to lightweight mirrors used in Synchrotron Light Telescopes and are located in the SPS and LEP rings and transfer channels. Most of these instruments are beam intercepting and therefore require very reliable drive electronics. As the SPS commenced operation in 1976 and LEP in 1989 they had completely different control systems. The coincidence of the upgrading of the SPS controls [1] in 1993, which necessitated the replacement of the stepping motor controls, and the planned energy upgrade in LEP, called the LEP 2 project, which required an additional 80 motorised collimator blocks, was the ideal opportunity to change to a common industrial position control system which would fulfil the requirements for both machines in the most economical way. Specific problems encountered in the SPS and LEP are the long distances between the motors and controllers, the high level of Electro Magnetic Interference (EMI) generated by the usual power electronics of industrial systems, enhanced by the long cable drives and the restricted space available for electronics in the LEP underground equipment caverns. The interface to the control system had to be able to accept different environments, currently PCs in the SPS and VME crates in LEP. The whole system had also to be reasonably modular so as to enhance maintainability and facilitate repairs by a single person during off-hours interventions.

2. SYSTEM REQUIREMENTS

The new system has to drive the existing motors and resolvers through the existing cables, which represent a major part of the cost of the motorisation. In the SPS, the stepping motors are of the unipolar type with six wires per motor, whereas in LEP they are of the bipolar type using four wires. In LEP the dynamic torque has to be above 1 Nm for driving heavy tungsten collimator blocks [2] or of 0.1 Nm for driving precision optical components [3], whereas in the SPS the required torque is 0.2 Nm for instruments located essentially in the transfer channels [4] between the PS, the SPS, the North and West Experimental Areas, and LEP. As most instruments to be driven in the SPS domain are in transfer channels, their positioning is less critical than in LEP where the majority of the controlled instruments are positioned close to the stored beam and require high precision and reliability to ensure long beam lifetime and equipment survival. For this reason an independent position measurement system for these instruments was justified. Resolvers were chosen for this task as their precision and reliability had already been demonstrated in previous projects [5]. Because of limited resources available in application software production, the new control system has to operate with existing application software, some programs being nearly twenty years old. The system has to be able to interface to two different environments, VME crates running under OS9 in LEP and PCs running LynxOS in SPS and possibly to other systems in the future.

The EMI noise generated by the system has to be kept to a strict minimum so as not to interfere with other accelerator components, particularly those distributed along the transfer channels and around the LEP ring, such as the high precision Beam Current Transformers, Secondary Emission Monitors, Beam Position Pick-Ups and the large experimental detectors in LEP.

Finally, the selected system has to be fully documented and preferably had to be a catalogue item. However, considering the specific demands, this condition was rather unlikely to be met, so the manufacturer had to guarantee to introduce the product into his standard product line for at least ten years, thus ensuring that CERN has the possibility of extending the system if needed in the future.

3. SYSTEM SPECIFICATION AND IMPLEMENTATION

The project was handled as a "farming-out" project with CERN producing a detailed functional specification and the hardware development being entirely the responsibility of industry. The most important parts of the specifications [6, 7] were the EMI requirements for the motor controllers and the system interface for both applications. The other characteristics such as dynamic motor torque and precision achieved with the resolver read-out are relatively standard. A complete system extended test was demanded in order to minimise the number of defects after reception at CERN.

3.1. EMI Specification and measurement

The EMI generated by the usual industrial stepping motor drivers is unacceptable in an accelerator environment. This became evident after an installation of industrial controllers in LEP [8] where filters had to be added to controllers, with pulse width modulation type drivers, because they generated unacceptable noise levels for other precision instruments [9]. This noise is in general not a problem in industrial applications where the cables between the motors and the drive electronics are kept short, and hence the radiated noise is acceptable. This is not the case in large accelerators and storage rings where the drive electronics has to be kept away from high radiation areas and is linked to the individual motors by long cables. Good experience had been gained with true DC powered drivers [5], where the noise is generated only during position changes and contains lower frequencies. A non-trivial problem was the definition of an "acceptable noise level". The availability of a recent Standard, CEI/IEC 478-5 [10], greatly helped in the specification of the acceptable noise. This standard is primarily intended for the measurement of the magnetic component of the local field of the DC output of stabilised power supplies, however it was found to be very useful for characterising the noise generated by stepping motor drivers. This standard uses a printed circuit board spiral antenna, Fig. 1, to measure fields in the frequency range of 10 kHz to 30 Mhz. This antenna was positioned at 1 cm from a flat cable inserted between the driver output and a 600 m cable linked to the motor and connected to a digital scope with an FFT facility. The noise level of the available filtered drivers was measured and used to define an acceptable noise power spectrum for the new system. The measured results from the original LEP drivers, with and without filter, and the new driver, with motor

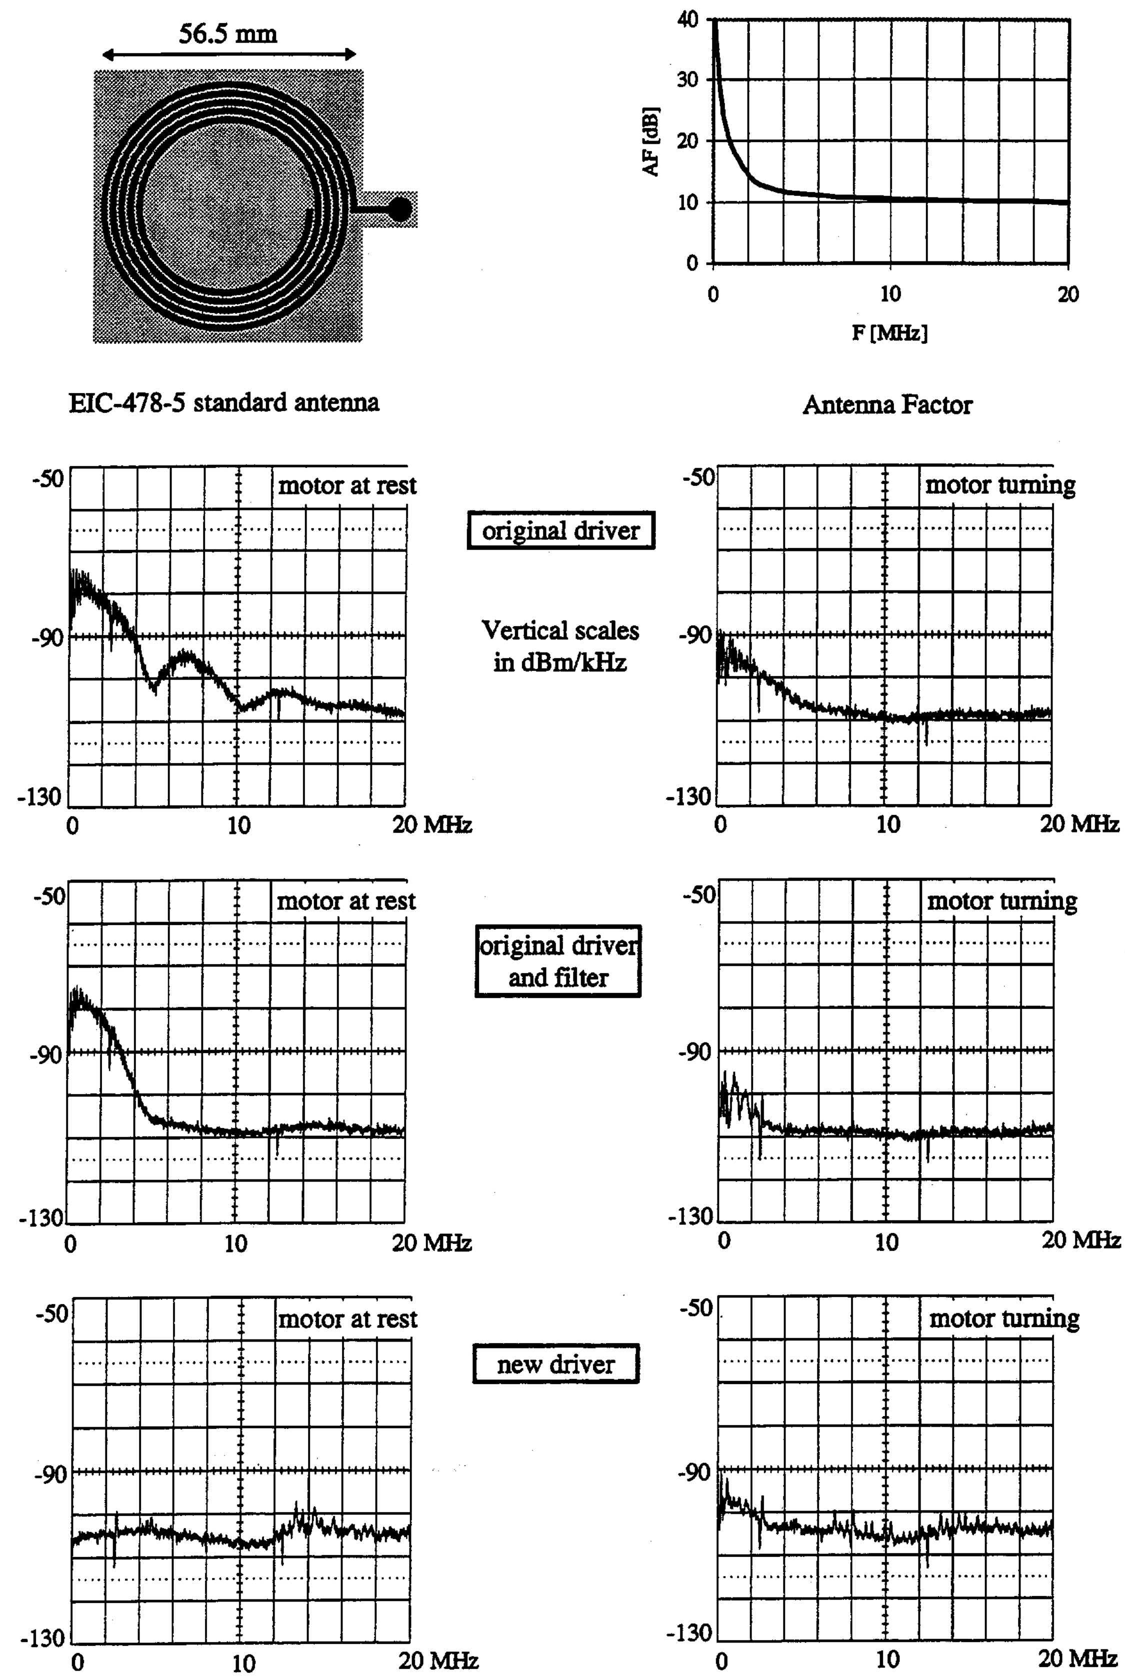

Fig. 1: EMI measurements.

at rest and in rotation are given in Fig.1. It can be seen that the noise generated by the new drivers with the motor at rest is very low, and that the noise with the motor in rotation stays below or close to the global envelope defined from the first generation filtered drivers. The Antenna Factor has to be added to these measurements to get real radiated EMI power [10].

3.2 Interface to the control system

The other important point was the interface to the two existing control systems and the possible use with other control systems, in particular the one which will be defined for LHC [1]. In all existing applications, the motor movements are infrequent and have to take place in time slots of the order of seconds at intervals of hours. The speed of the motor is matched to the task and load, and is in general 5 turns/s. In some applications, the motor may have to be stopped quickly to protect a circulating beam. A serial communication link to a cluster of motor controllers is adequate for fulfilling these requirements. An RS 232 link was specified as the motor controllers can always be located close to such a communication port. After discussion with the manufacturer [11] of the chosen system, an RS 485 link was also made available which will allow in the future the use of motor control clusters located far from a communication port. It was specified that one communication port should be able to control up to 64 motors and readout 64 resolvers. These numbers are large enough to cover all the needs in both the auxiliary buildings of the SPS and the underground caverns in LEP in the foreseeable future. Because of the universal nature of this port, the control and readout system can be connected to any control system. The present structure is depicted in Fig. 2, where both situations in SPS and LEP are represented on the same Ethernet ring. The LEP resolvers are interfaced in the same way as the motors. This architecture is economical as it doesn't need dedicated VME crates with their cost overheads and limited capacity and can be expanded easily to any number of motors.

The controls were specified to be a simple command/response dialogue for robustness. An independent "Reset" facility was demanded, this being necessary for restarting the system after a major power line perturbation. This feature is particularly useful in LEP where the underground equipment caverns are located several kilometres away from the Control Room and the Office and Laboratory Buildings.

4. HARDWARE DESCRIPTION

4.1 Motor Controllers

The system is based on an existing six-high double width control card for driving four stepping motors. This card, which can drive variable reluctance and hybrid stepping motors, in unipolar or bipolar mode, defines the basic modularity of the system. It is controlled by a serial link of type RS 232 or RS 485 and has one microcontroller per motor, with control and communications libraries available. Eight logic inputs per motor, which can be treated within 1 μs, can be used to initiate real time sequences. The motor control, the eight logic inputs and the serial link are treated simultaneously. The system accepts two end switches which are monitored continuously. The power is generated in one unit common to eight motors. The motor phases are driven by individual current generators, and the currents can be adjusted between 0 and 4 A per motor, with a maximum voltage swing of 96 V. The stand-by current is adjusted separately from the dynamic current, this minimises the power consumption and EMI generated at rest. The phase currents in each motor are generated by switching-type power supplies with filtering at the source and are controlled by H-type bridges using MOSFETS. Utilising only one current generator for the two motor windings results in a smoother waveform therefore minimising mechanical vibrations and generating smoother motor rotation. This also decreases the energy of the radiated higher harmonics. Taking into account the mode of operation of the motor system, the total installed power was minimised by limiting the simultaneous use of the most powerful motors to one motor per control card, i.e. 1 out of 4. Motors at a distance up to 1500 m from the driver can be controlled reliably. The calculated MTBF of an eight-motor controller assembly is greater than 10^5 hours.

4.2 Resolver Read-out System

The system is modular and is built up using six high single width eurocards, each one operating on eight resolvers. Each card is controlled by a microcontroller of the same type as used for the motor system. Its task is to control the input 8-channel multiplexer, the 14 bit Synchro-to-Digital converter and the resolver

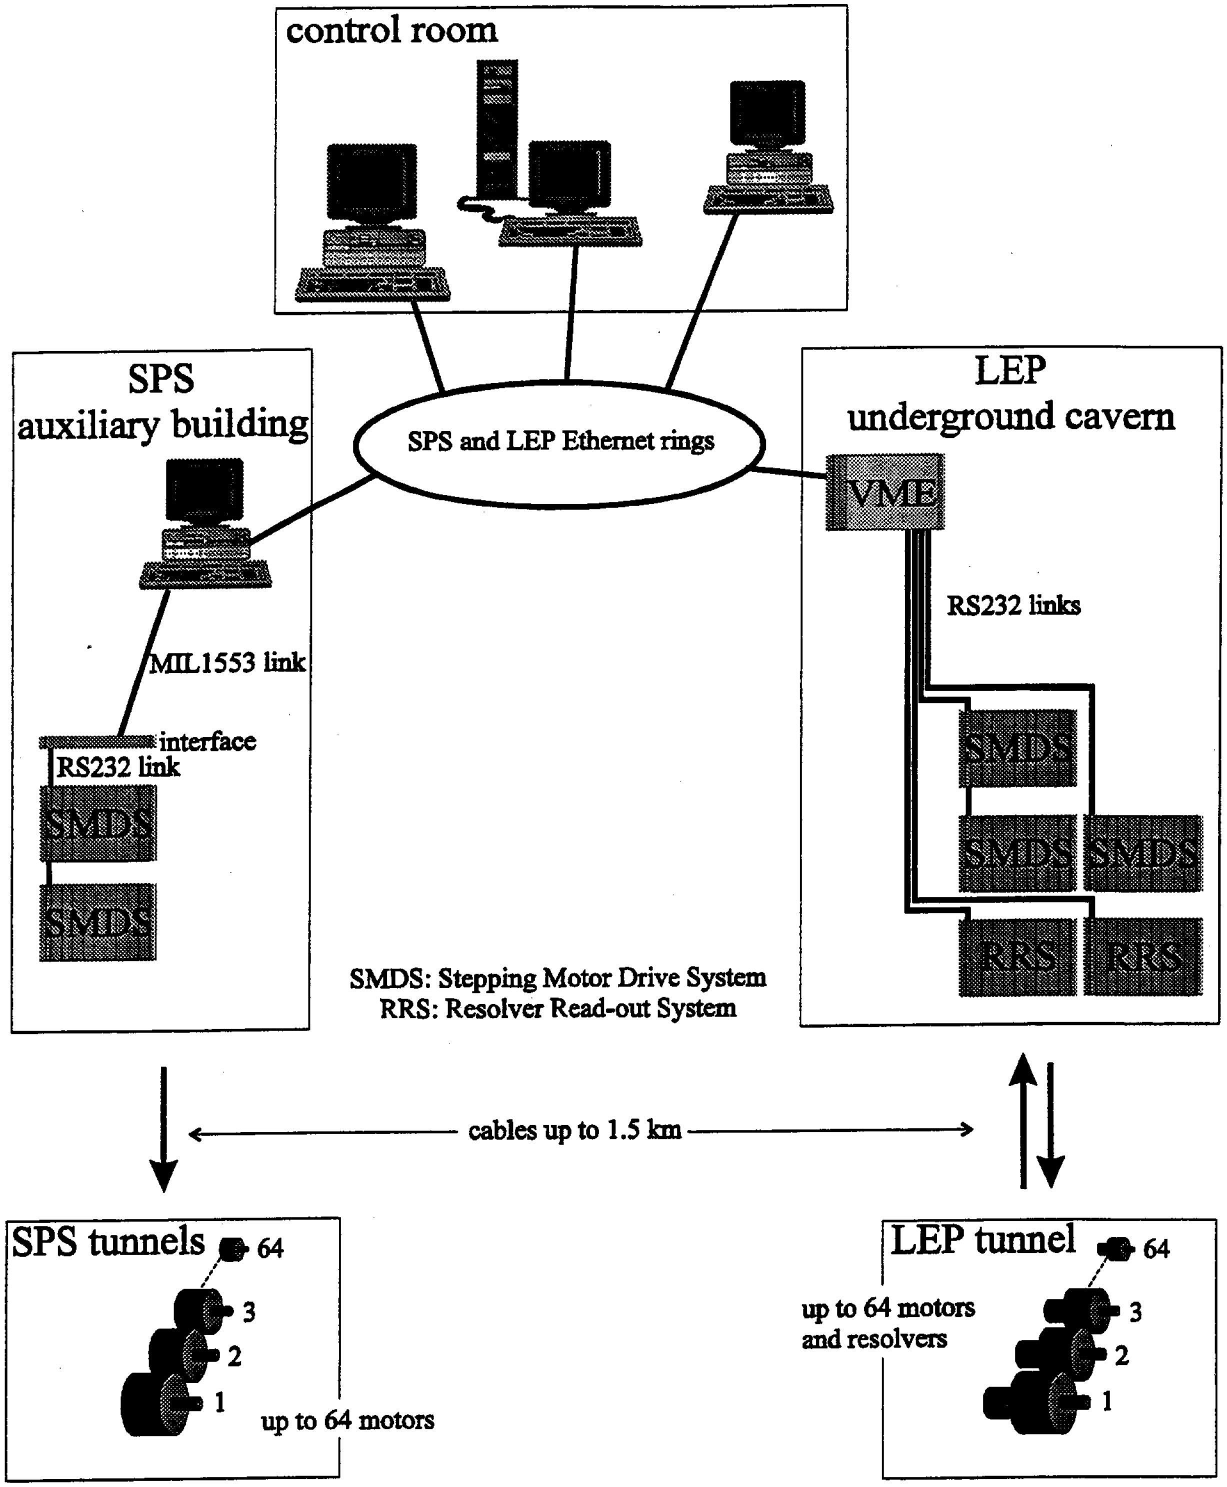

Fig. 2: Architecture of the SPS (PCs) and LEP (VME) motor control system.

excitation level. It performs calculations on the raw data and has to handle the communications over the serial port. The resolver excitation level is kept constant irrespective of the distance to the resolver, by checking the level of the combined return signals and controlling programmable gain amplifiers. The input signals are scanned in turn using a differential multiplexer followed by two instrumentation amplifiers for impedance matching and common mode rejection. It takes 1s to readout the positions of the eight resolvers connected to the card. The serial link can read up to 99 resolver cards, and uses the same protocol as for the motor control system, thus allowing the use of the same output port for controlling motors and reading resolvers. Separation of the motor and resolver functions are made by the component address, the motors being within the first 64 locations and the resolvers in the last ones. In practice this facility is not used to its maximum capacity because it limits the speed of execution. Average and RMS calculations are made on the raw data for better precision and measurement quality evaluation. The direction of rotation and the mechanical offset are programmable for each resolver. This feature eases the mechanical set-up and improves the absolute precision of the measurement. The calculated MTBF for an eight-channel readout module is also greater than 10^5 hours.

4.3 Layout and installation

The motor and resolver systems were specified to be reasonably modular in order to achieve a good compromise between the cost of an installation, determined also by the unused installed capacity and the cost of the spare elements which have to be distributed over the site for efficient off-hours interventions, and its compactness. The basic modularity is of eight motors or resolvers, with a sub-modularity of four motors of the same type. Each Power Unit is connected to a protected 220 V mains outlet. The system uses the six-high euro-chassis standard. The components are also designed so as to have the possibility to have eight motor controls and eight resolver readouts housed, together with their RS 232 port, in one single chassis: Fig. 3. This gives enough flexibility to build up a motor control station in the most compact and economical way. There is in fact little unused capacity in the present installation. As requested, the system does not need any forced ventilation nor other specific cooling. Taking into account previous experience, a lot of consideration was given to the connection of the system to the cables going to the elements in the tunnel. The original LEP system used a number of cable patch panels and jumpers which, apart from their cost, complicated the installation and the trouble shooting and were not favourable for the overall system reliability. The new installations have no patch panels, the cables arriving from the tunnel are fixed on staggered bars and connected to the motor and resolver chassis by one short cable per motor or resolver. The installation, as can be seen in Fig. 4 showing racks for 28 motors and 28 resolvers, uses less space, is more economical and easier for maintenance and trouble shooting than before. Up to 48 motor controllers together with 64 resolver read-outs could be installed in one rack. The installations have at present been limited to 32 motor and 32 resolver interfaces per rack, in order to have all elements for the controlled instruments in the same rack and still have an easy access to the individual cables leaving the racks for the tunnels. The RS 232 links go to PCs or VME crates located in other parts of the same equipment building.

5. SOFTWARE IMPLEMENTATION

The interface of the RS 232 link available with the motor controllers and resolver readouts can activate all the basic controls of the motor drivers and resolver digitisers. However, the system has also to be able to interface to two different environments, be compatible with the existing application software and deal with two different ways of operating the motors.

In LEP, on the one hand, the existing operational interface has to be conserved because the original motors are still to be controlled by VME crates running under RMS68K, and on the other hand existing VME crates running OS9 are to be used for economy reasons to install the RS 232 ports for controlling the LEP 2 motors. A module with four independent RS 232 channels is used to increase the throughput to the motor controller system which is limited in speed to 4800 baud per RS 232 input. A method able to handle the 172 collimator motors as a single instrument had to be defined to control all the motors in a single command within a reasonable time. The software structure arrived at is outlined in Fig. 5. The 'communication' process handles all the requests from the client applications and dispatches them via shared memory to the 'control' tasks. Each 'control' task handles the requests to the motors linked to its RS 232 line.

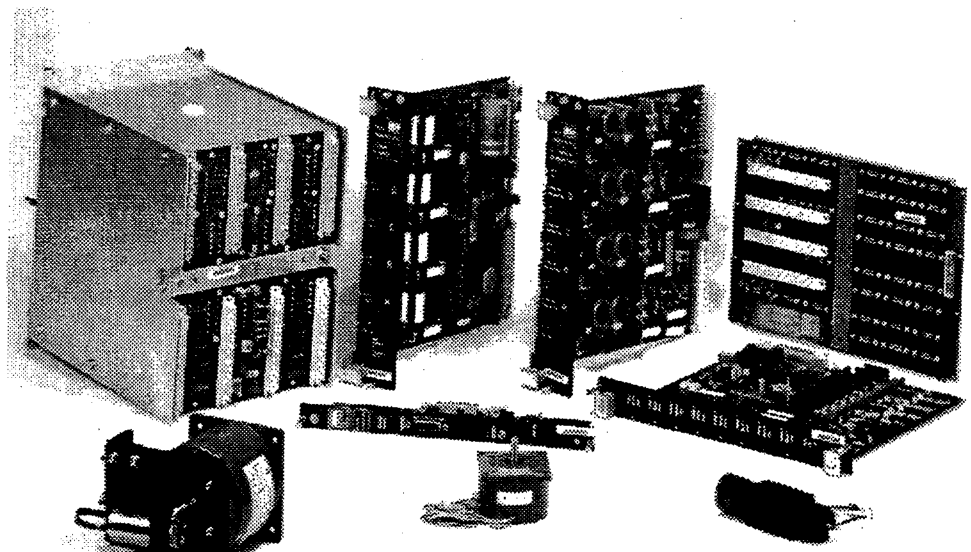

Fig. 3: The elements of SPS and LEP position controls: from left to right: at the rear: Power unit for 8 motors, intelligent 4 motor controller and 4 motor current generator, 32 resolver backplane, in the middle: RS 232 / RS 485 interface and 8 resolver readout unit, in the front: 1 Nm motor and resolver, 0.2 Nm and 0.1 Nm motors used in LEP and SPS.

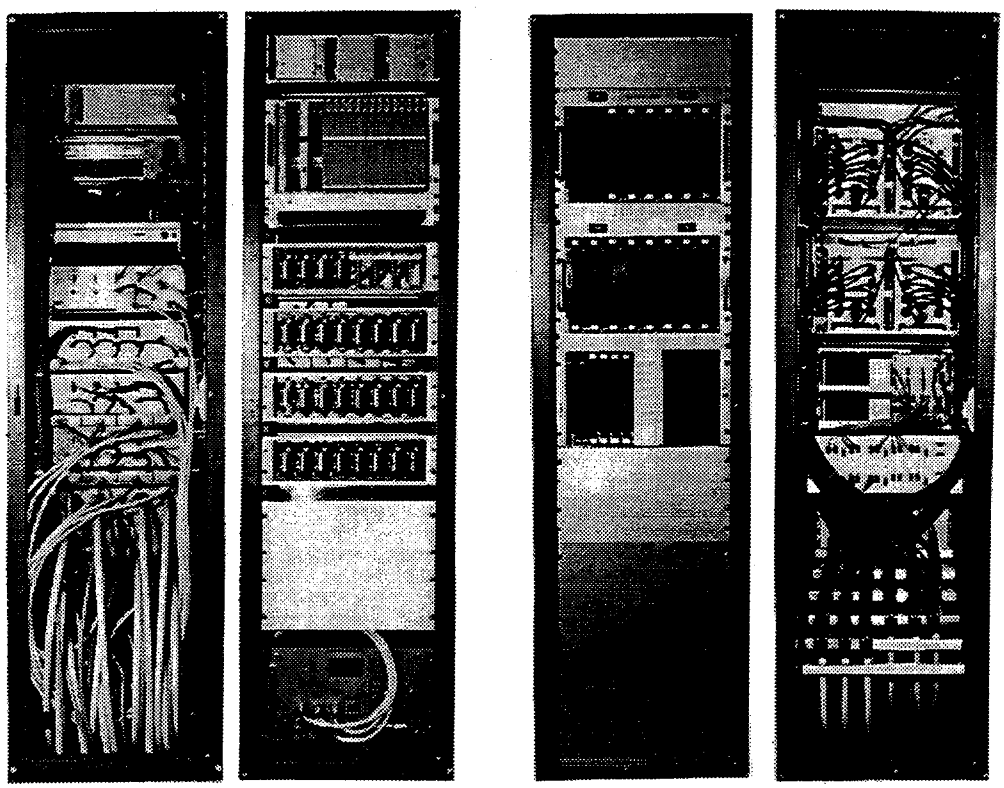

Fig.4: Rear and front views of the racks with control electronics and cable connections for 28 motors and 28 resolvers with the original LEP (left) and with the new (right) systems.

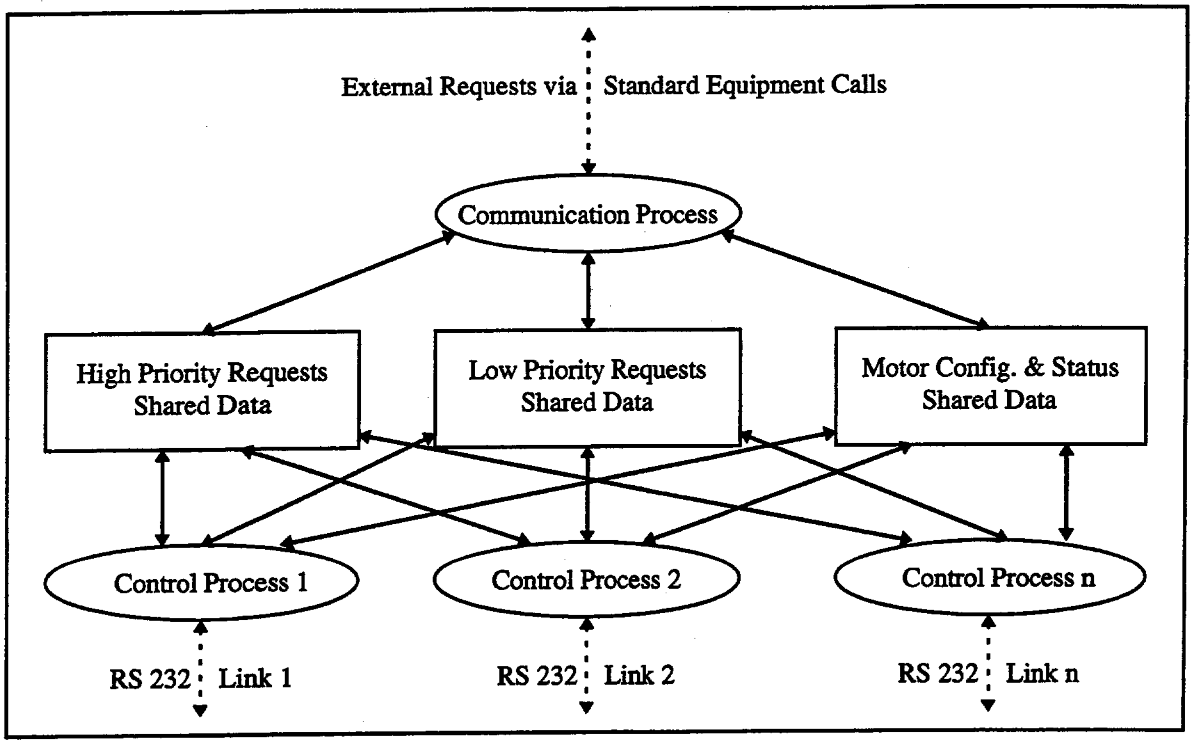

Fig. 5: Low Level LEP Software Architecture.

The software has to implement two different levels of priority for the client requests. Some requests, like the 'stop motor' command, have to be executed immediately and are called 'High Priority Requests', other commands can be delayed and are referred to as 'Low Priority Requests'. This situation is handled with the help of Shared Memory blocks which are available to all processes and contain structures describing the requests, the arguments and responses for each motor. The 'Motor Configuration and Status' Shared Memory contains structures to describe the motor settings (speed, position limits) and status (position, limit switches).

When the 'communication' process receives a status request, it reads the corresponding Shared Memory block and sends the answer back immediately. If it receives a High Priority Request, it stores it in the proper block, waits for the 'control' process answer and sends the latter back immediately to the client. When a Low Priority Request is received, it is stored in the corresponding memory and the client is informed that a request has been recorded and should be treated as soon as possible. The client has then to check if the action has been completed. The 'control' processor repeats indefinitely for each motor the following loop: execute any pending High Priority Request - execute motor 'n' Low priority Request or read its position - and so on, scanning all motors. With this procedure a global status or Low Priority command over the 172 collimator motors is executed within five seconds and a High Priority command is handled in one second. This is quite sufficient for all remote control applications. Actions with shorter reaction times, necessary for treating end switches, are dealt with locally in the motor controller using the logic inputs and take a few microseconds. It has been very easy to incorporate additional motors and resolvers into the system by using this protocol.

The situation is simpler in the SPS, each motor being treated as a single piece of hardware because the motors belong to different kinds of instruments. However a special procedure had to be developed to access the equipment via PCs running LynxOS. This procedure has to transform the application requests into RS 232 command/response strings using existing equipment access paths. The main task was to deal with the 'standard equipment access' [1], to translate the user requests and send them through the existing MIL 1553 Bus: see Fig. 2. This was simplified by using a 'black box' [12] which has been developed to translate the MIL 1553 message into an RS 232 message and vice-versa. New features, like the remote reset and detailed status reports, which were not available in the past have been implemented, thus improving significantly the diagnostics and trouble shooting on the equipment.

6. TIME SCALE OF THE PROJECT, PERFORMANCE AND PRESENT EXPERIENCE

6.1 Time scale of the project

The specification for the motor controllers was finalised in February 1993, the call for tenders was sent out at the beginning of March and the analysis of various proposed systems started in May. The first selection criterion was the EMI noise spectrum. All but one system were eliminated at this stage. Detailed discussions started then with the remaining company to match CERN's requirements and wishes with the company's perception of future market needs. This phase was extremely positive for both parties and in August 1993 an order was placed for the controls for 340 motors. The various components of the system were delivered between November 1993 and March 1994 and 210 motors, i.e. all the SPS and the then installed LEP 2 motors were equipped with the new control system for the 1994 start-up. The system was up and running just one year after the call for tenders had been made.

The resolver project was launched in August 1993 and handled in a similar way to the motor driver project. The best proposal was made by the same company who made the motor drivers and subsequently an order was placed with them. The elements for reading 136 resolvers were received before March 1994 and were installed ready to read the positions of the LEP 2 collimators for the 1994 start-up.

The prolonged tests at the manufacturer's premises resulted in practically no defects after reception of the components at CERN and helped greatly in keeping to the tight schedule.

Additional modules for equipping the 120 original LEP motors have now been ordered and will be installed during the 1995/96 winter shut down.

6.2 Performance and present experience

For the LEP motors, a higher torque, 1.2 Nm against 1 Nm, was obtained safely with the new system, this permitted the use of the same motor type for all heavy load applications in LEP, instead of two motor types as had previously been used. The noise level generated by the stepping motor drivers and seen by the other instruments is now below their sensitivity level to external perturbations. Moreover, other noise sources in the accelerators have now become the dominant perturbation [9]. The true DC currents for the holding torque when at rest, which is for most of the time, is particularly beneficial from the noise point of view. The automatic excitation level adjustment of the resolvers keeps the signals constant irrespective of the distance to the controlled instrument and guarantees a nearly constant position readout precision of 13 bits, i.e. 8192 counts. The precision of the position readout is now limited by the gear coupling the resolver to the motor, so a new backlash-free precision worm wheel assembly has now been implemented [13] which more closely matches the achievable readout precision. As the mechanical coupling error curve is smooth and stable, a simple harmonic correction can be introduced in critical applications to bring the precision of the position readout down to twice the resolution of the motor in half-step mode, i.e. 10 μm for an 80 mm stroke, which is at least five times better than needed in standard applications.

No major problems were encountered in the SPS despite the large number of motor controllers changed at the same time. The new motor control system required no effort to be invested in applications software changes. In both machines, the change of control electronics was transparent to the users, which is the best sign for a successful change. No faults have been experienced between March 1993 and October 1995 with the new control electronics driving the 210 motors and 80 resolvers, distributed over 14 buildings and underground areas. It can therefore be claimed that the system is extremely reliable.

7. CONCLUSION

The project of the SPS and LEP 2 motor control was handled as a "farming-out" project. It is considered successful because it was possible to obtain in a rather short time a good technical solution, which was economical and profitable to both CERN and the industrial company concerned. To achieve this success, the combination of several factors was necessary. On the client side (CERN), expertise was necessary in order to define the boundary conditions for the equipment: a good understanding of stepping motors and resolvers; the minimum performance needed, the precision achievable with resolvers, the interface to the control systems, the understanding and precise measurement of the EMI noise, flexibility to consider the industrial partner's priorities, i.e. the potential market for the product. On the industrial company's side, technical

expertise, good experience with standard industrial applications, flexibility to understand the client's needs and willingness to match them with the demands of potential markets including the scientific community at large, were necessary. This left CERN staff, after an initial phase of detailed specification production and system selection, time to consider in depth specific problems, like rack layout, cable interface and diagnostic software, which are important in the long run for good maintenance and availability of the equipment and are often neglected because of lack of time. Added advantages for CERN are the now common expertise of the original two SPS and LEP teams and the reduced common spares stock for the two accelerators.
The overall experience is positive for both parties.

Acknowledgements
C. Bovet and K.H. Kissler supported the project. Interesting discussions took place with G. Baribaud about system interface and with R. Rausch for the SPS integration. L. Symons introduced us to Standard CEI/EIC 478. P. Charrue was instrumental in the success of the integration of the system in the SPS/LEP controls. B. Denis integrated the new collimator controls in the LEP collimator application program. G. Burtin and J.P. Bindi took over the LEP installation and all mechanical improvements and performance measurements and J.C. Lucas tested all SPS and LEP modules before installation and implemented the SPS installations. Their contributions are acknowledged with gratitude.

REFERENCES

[1] R. Lauckner, R. Rausch: SPS and LEP Controls, Status and Evolution Towards the LHC Era, this Conference

[2] R. Jung et al.: Design of a new generation of collimators for LEP 200, Proc. of the 1993 IEEE Part. Acc. Conf., Washington D.C., 1993, and in CERN SL/93-20 (DI), May 1993

[3] C. Bovet et al.: The LEP Synchrotron Light Telescopes, Proc. of the 1991 Part. Acc. Conf., San Francisco, 1991, and CERN SL/91-25 (BI), April 1991

[4] J. Bosser et al.: Beam transfer monitors for the operation of the SPS with oxygen, Proc. of the 1987 Part. Acc. Conf., Washington D.C., 1987, and CERN/SPS/87-9 (ABM), February 1987

[5] R.J. Colchester et al.: Position control in radiation environments, experience from the ISR, Possibilities for LEP, CERN/LEP-BI/86-08, February 1986

[6] Supply of a Stepping Motor Drive System, Specification IT-2 158/SL, February 1993

[7] Supply of a Resolver Readout System, Specification CO-13 832/SL, August 1993

[8] C. Beugnet, K.D. Lohmann: Private communication

[9] B. Dehning, P. Puzo: Private communication

[10] International Standard CEI/IEC 478-5: Stabilised power supplies, d.c. output Part 5: Measurement of the magnetic component of the reactive near field, 1993

[11] Midi Ingenierie, Labège Innopole, BP 131, F 31 676 Labège, France

[12] P. Charrue, M. Clayton: The New Control Infrastructure for the SPS, this Conference

[13] G. Burtin: Private communication

REMOTE CONTROL OF A STREAK CAMERA FOR REAL TIME BUNCH SIZE MEASUREMENTS IN LEP

A. J. Burns, J. J. Savioz, J. C. de Vries
CERN, SL Division, CH-1211 Geneva 23

Abstract

A double sweep streak camera, built by industry according to CERN specifications, has been used for a number of years to provide real time three-dimensional measurements of bunches in LEP, by means of a dedicated synchrotron light source. Originally requiring local manipulation in an underground lab close to the LEP tunnel, the camera can now be fully operated via the control system network. Control functions such as the adjustment of lens and mirror positions, the selection of camera sweep speeds and the setting of 12 ps resolution trigger timing, are handled by various networked VME systems, as is real time image processing. Bunch dimension averages are transferred every few seconds via the control system to the LEP measurement database, and a dedicated high bandwidth video transmission allows the streak camera images and processed results to be viewed in real time (at 25 Hz) in the LEP control room. Feedback control loops for light intensity, trigger timing and image tracking allow the setup to provide useful bunch images and logged measurements over extended periods, without human intervention. An X-Window based control application (GUI) allows LEP machine operators to select different bunches for display and measurement. The same application permits the specialists to control all parameters of the system.

1. INTRODUCTION

Synchrotron light pulses are produced by the passage of e^+ and e^- bunches through small wiggler magnets placed at 67 m on either side of intersection point 1 of LEP [1]. The light is extracted by two thin beryllium mirrors in the vacuum chamber and focused via two evacuated optical lines on a double sweep streak camera [2] in an underground optical laboratory 15 m from the tunnel [3]. The optical setup allows the simultaneous observation of the side and top views of any photon bunch from both LEP beams within the same fast sweep. Measurements can therefore be made of the instabilities and sizes of the particle bunches in all three dimensions [4]. The photon bunch length and longitudinal density distribution corresponds to that of the particle bunch that emitted it. The absolute calibration of the transverse measurements is more difficult, being influenced by the optics of the source and the light collection and transport, and is better done by other instruments that are dedicated to emittance measurements [5]. The slow sweep allows up to 100 fast sweeps to be recorded on one image, which can be used to follow successive bunch passages.

The streak camera system can be decomposed into a number of major blocks requiring control, as shown schematically in Figure 1. First, there is the synchrotron light path that brings light pulses from the LEP tunnel to the entrance of the streak camera. This includes a light intensity feedback system to maintain the light incident on the streak camera at an optimum level, as well as a number of mechanically controlled lenses, mirrors, shutters, and camera rotary switches. This subsystem is described in Section 2. Next, there is the synchronisation between the beam signals, the streak camera itself, its associated CCD camera and the image processing system. This covers the range from milliseconds down to picoseconds and is described in detail in Section 3. Some features of the image processing system, based on the Datacube MaxVideo20 board [6], are described in Section 4. Finally, aspects of the control software, including the X-Window control application being prepared for this entire system, are covered in Section 5.

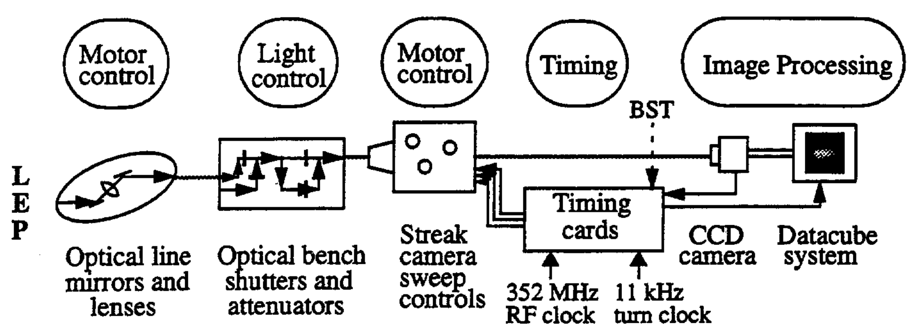

Figure 1 : Principal streak camera subsystems

2. The Synchrotron Light Path

2.1 Lenses and mirrors in optical lines

The optical lines that transport the synchrotron light from each beam to the optical laboratory are each equipped with an achromatic lens that can be moved longitudinally over a range of 75 cm in order to bring the light to a focus on the streak camera at a distance of about 20 m. These movements are motorised via stepping motor interfaces standardised for LEP beam instrumentation [7], and resolvers provide measurements of the achromat positions. Also motorised are the inclinations in two orthogonal planes of various mirrors used to centre the light beam in the optical lines, including the beryllium mirrors used to extract the light from the LEP vacuum chamber.

2.2 Synchrotron light intensity control

A feedback loop maintains the synchrotron light intensity received by the streak camera within a limited range. The aim is to have as bright as possible an image on the CCD camera without saturation, thus optimising the signal to noise ratio. The optimal incident light level of course depends on the streak camera sweep settings. Slower sweeps concentrate the bunch image on fewer CCD pixels and require more attenuation. Taking this into account ensures that there is no risk of damaging the streak camera phosphor screen by concentrating too much light on it.

A fixed small proportion (4%) of the incident synchrotron light is deflected towards photomultipliers (PMs), one for each beam. Analogue electronic circuitry, including a peak detector and fast comparator, processes the PM pulses generated by the individual LEP bunches, producing a level that is proportional to the highest pulse height in the last few tens of milliseconds. This level is digitised by an 11 bit ADC every 100 ms; availability of a value triggers an interrupt service routine (ISR). If the measured signal level falls outside the presently required range, then the ISR calculates the change in the attenuation required to re-establish the central value of the range.

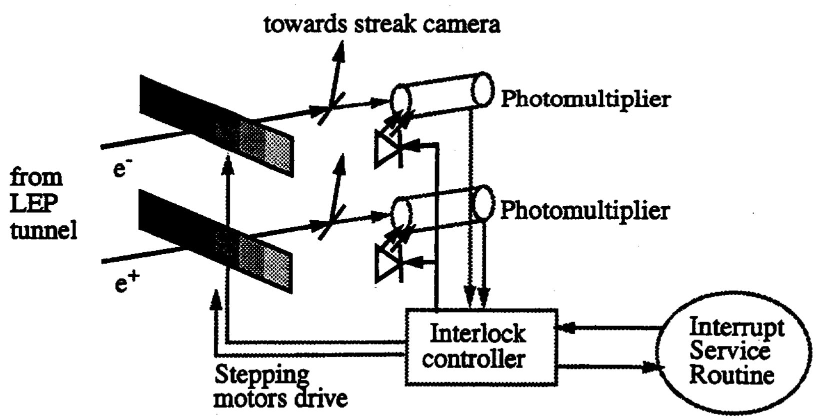

Figure 2 : Synchrotron light intensity control

Optical attenuators, 150 mm long, are mounted on sliders driven by stepping motors at a speed of 40 mm/s, with a step size of 63 μm. The light attenuation obtained increases exponentially with distance up to a maximum value of 1000 (i.e. from 100% to 0.1% of incident light transmitted). The 4 mm movement that can be produced in the 100 ms between interrupts therefore corresponds to a 20% change in attenuation. The ISR derives the number of motor steps required, assuming that the attenuation changes linearly within the 4 mm interval. The obtained value is limited to 64, the maximum number of steps that can be executed in 100 ms, the excess being treated in subsequent intervals. To compensate for variations in the PM gains and ensure that the reference signal level always corresponds to the same incident light level, a calibration procedure is used to adjust the PM bias voltages to obtain the required reference level from a light emitting diode (LED) included in the PM bases.

In addition to the software-controlled light attenuation feedback, which can take up to 4 seconds to reach maximum attenuation, there is a fast analogue detection of preset maximum and minimum light levels. The maximum level is particularly important as it has been determined to ensure that the photocathode of the streak camera is not damaged by very intense light pulses. As the streak camera receives synchrotron light from a region of LEP where an intense laser beam is brought into collision with the particle beam in order to measure polarisation, there is always a possibility that a reflection from this laser is directed towards the streak camera. The analogue detection circuitry generates an interlock that closes a shutter in front of the streak camera when the maximum light level is exceeded for more than 20 ms. In addition, an error condition detected on either PM high voltage supply generates an interlock. However, in the case of normal operation of the polarimeter this level of protection should not be necessary, as there is also a hardware veto signal generated by switching on the laser that produces an interlock. In addition, there is a software call made to the low level streak camera light control software that closes down the streak camera system elegantly when the LEP operators request that the polarimeter laser power supply be turned on. This shutdown consists of closing the main shutter, disabling the light intensity control system and setting the attenuation to maximum. Once the laser has been turned off another call generates a warm start of the system. The light intensity control system is first calibrated and then enabled. Once it is stable the main shutter is reopened. Even once the conditions that caused it have disappeared, the interlock always remains active until reset remotely by software.

A VME module called the Interlock Controller [8] has been developed for measuring the pulses from the PMs and controlling their bias voltages and test LEDs, controlling the attenuators via stepping motors, detecting the high and low light levels, handling the interlock and operating the view shutters described below.

2.3 Particle beam and view shutters on optical bench

Before being combined into one beam on the optical bench, the light emitted by the e^+ and e^- beams passes through electro-mechanical shutters that can be remote-controlled via the Interlock Controller. The combined e^+/e^- light beam is then split into two paths, one containing a Dove prism that rotates the photon bunches by 90° about their longitudinal axis, the other containing a manually variable optical delay that provides a time separation between the rotated and unrotated bunches after recombination in front of the streak camera. On the final image, one light path provides the view of bunches seen from the side, while the other provides the top view. Shutters in the two paths allow the selection of either or both views. The time difference of a few 100 ps introduced by the optical delay and the fixed 500 ps e^+/e^- difference makes it possible to display both top and side views of both beams on the same image. In the case where it is required to change from one view of different bunches of the same beam (such as shown in Figure 7) to another beam/view combination of the same bunches, the control software is able to use the known time separations to adjust the fast trigger delays to keep the bunch images in the same positions on the screen.

2.4 Streak camera sweep controls

At the time of construction of the streak camera it was not considered necessary to have the possibility of controling the sweep speeds remotely. A special mechanical assembly therefore had to be constructed later to allow the 3 main control rotary switches to be turned remotely by stepping motors. These motors are controlled by the same stepping motor interfaces used for the mirror and lenses of the optical lines [7]. The software must keep track of the rotary switch settings as there is no measurement available of the motor positions. There are 2 discrete position rotary switches, one for the fast sweep settings of 0, 90, 150, 200, 300 and 1200 ps/mm (measured values) and the other for the slow sweep settings of 0, 4, 40 and 400 μs/mm. The sweep values correspond to displacements measured on the image incident on the CCD chip, which has a pixel size of 23 μm. In addition there is a continuous slow sweep control rotary switch, which in 10 revolutions varies the slow sweep from any set non-zero value to that 10 times slower (e.g. from 4 to 40 μs/mm). Is it thus possible to obtain any value from 4 μs/mm to 4 ms/mm, corresponding to a time range on the 288 pixels of the CCD in the direction of the slow sweep from 26 μs to 26 ms (horizontal in streak camera but vertical on final image -- see Figure 7). However in practice the largest useful time is limited to about 10 ms, as discussed in Section 3.1. The CCD chip has 384 useful pixels in the direction of the fast sweep resulting in a set of effective screen widths of 0.8, 1.3, 1.75, 2.6 and 10.6 ns. Changes of sweep frequency must be communicated by the control software to the light intensity control system, to the timing control software and to the image processing software.

3. SYNCHRONISATION

Successful operation of the streak camera system clearly requires synchronisation between the different entities involved, i.e. the streak camera itself, the CCD camera that records the image produced on the streak camera phosphor screen after passage through a multi channel plate (MCP) image intensifier and finally the commercial image acquisition and processing system from Datacube that digitizes the CCIR composite video signal from the CCD camera (of frame frequency 50 Hz, line frequency 15.625 kHz, and pixel frequency 7.375 MHz). In addition one has to synchronise with the synchrotron light pulses being received from the two circulating beams in the LEP machine. However it should be noted that light pulses are available from the beam at a high frequency (11.2 kHz for any individual circulating bunch) and thus in general only the streak camera needs to be precisely synchronised with the LEP beams. A special mode of operation, referred to as the "BST mode", is under development for synchronising bunch visualisation with particular machine events such as the instant of injection. This mode, which also allows the variation of the image frequency and the interval between visualised turns, is described in Section 3.3 below.

3.1 CCD mode

In the normal mode of operation, referred to as "CCD mode", it is the CCD camera 50 Hz frame clock that controls the synchronisation. As shown in Figure 3, it is the "CCD ready" signal produced by the CCD controller a few ms after the start of each alternate frame that launches the sequence of streak camera sweeps and also triggers the Datacube acquisition. The image from the streak camera is integrated on the image zone of a Thomson TH 7863 frame transfer CCD chip in the milliseconds following the "CCD ready" pulse in one 20 ms frame, and is available to the Datacube in the following frame of the composite video output.

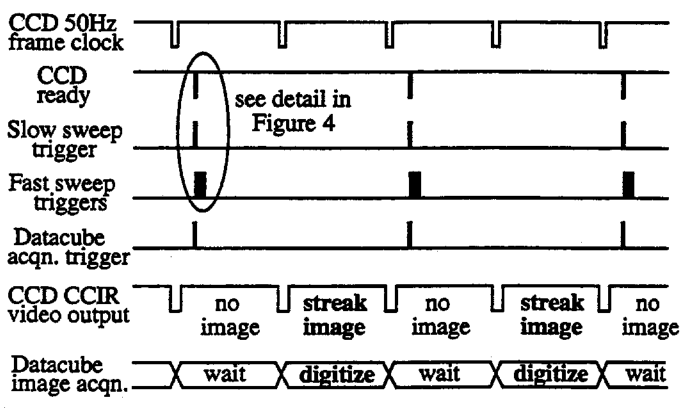

Figure 3 : Synchronisation of streak camera, CCD camera and Datacube system

The integration of the light from the entire burst of bunch images must be complete before the end of the frame and the transfer of the image into the memory zone of the CCD, from which the charges are read out. Taking into account the relatively slow decay time of the P20 phosphor at the exit of the MCP intensifier, the burst cannot cover more than about 100 LEP turns. The Datacube MaxVideo20 module operates in a continuous loop in which once it has completed digitizing one video frame it waits for the arrival of the trigger before starting the digitization of the next complete frame. This triggering scheme thus provides a new image for display and analysis every 40 ms.

The synchronisation of the streak camera with the synchrotron light pulses produced by the e^+ and e^- bunches circulating in LEP is done in various stages. The streak camera is equipped with a fast internal pulse bias of the photocathode that acts as an electro-optical gate. This gate must be triggered once for each LEP bunch passage to be measured. The time the photocathode remains sensitive can be varied between 0.2 and 2 µs, but is normally set to the maximum value for convenience. It is clear that to measure a given bunch on successive revolutions one needs to generate a trigger with a fixed delay with respect to a clock train synchronous with the circulating beam. This 11.25 kHz clock, referred to as the "LEP turn clock", is derived from the LEP RF system and transmitted via optical fibres to the underground laboratory containing the streak camera. It has an rms jitter of about 1 ns. A general purpose delay module (LSD [9]) is used to produce pulses with a programmable delay and width in 50 ns steps with respect to the turn clock.

The bunch structure of the LEP beams in 1995 consisted of 4 equidistant bunch trains each containing 1-4 bunches separated by 87 RF buckets (87 τ_{RF} = 247 ns). One LSD output is used to produce the triggers for the optical gate, allowing the selection of 1-4 bunch trains per turn. As the full bunch train passes within 750 ns, all bunches in a given train are accepted by the optical gate width of 2 µs. It is the fast sweep triggers, generated with a special picosecond timing module (see Section 3.2) that select the particular bunches within a train and also determine their position on the phosphor screen in the direction of the fast sweep. The synchrotron light pulses from the other bunches that fall on the photocathode when the optical gate is open do not fall on the visible part of the phosphor screen. The streak camera slow sweep is used to separate the bunch images from successive LEP turns on the phosphor screen. To optimize the linearity, the initial and final 8% of the slow sweep ramp is off the visible part of the phosphor screen and therefore the slow sweep trigger must precede the first optical gate trigger by about 10% of the useful sweep time, a delay that varies with the slow sweep speed selected. A second LSD output is therefore used to generate one slow sweep pulse per turn that can be positioned to include the part of this relative delay corresponding to the fractional part of a LEP turn. In practice this level of refinement is not exploited, although it could be used to ensure that the first bunch image on the screen is optimally positioned.

The specific streak camera "Timing Adaptor" [10], takes as input the trigger pulses already correctly positioned with respect to the LEP turn clock and generates the burst of triggers for each full image. In particular it creates the integer part of the slow sweep to optical gate delay, by delaying the optical gate triggers by a fixed 8 LEP turns and the slow sweep trigger by a variable 0-7 turns. As shown in Figure 4, which gives as an example the case of 4 bunches acquired for 2 consecutive turns, these delays apply from the turn clock pulse received immediately following the "CCD ready". The variable slow sweep delay, as well as the number of LEP turns in the burst, are presently set manually on the front panel of the model, but will be programmable in the final version.

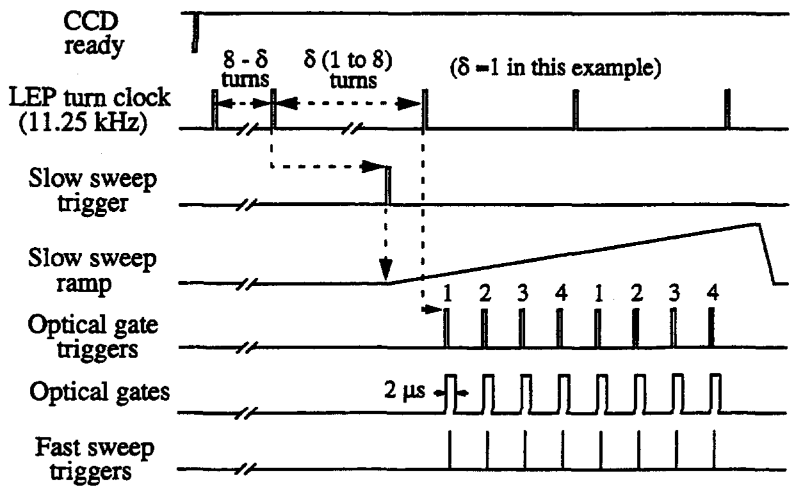

Figure 4 : Example of generation of triggers to measure 4 bunches on 2 consecutive LEP turns

The Timing Adaptor, also applies two constraints on the streak camera triggering necessary for the safe operation of the camera. The first is that the minimum time between triggers be 20 µs. In practice this fits well with the 22 µs between the bunch trains in LEP and simply limits the system to acquiring 1 bunch per train. The second constraint is on the average pulsing rate which should not exceed 200 every 40 ms. This limits one CCD image to containing, for example, 4 bunches per turn for 50 turns.

3.2 Picosecond timing

A picosecond timing module has been developed to provide beam synchronous trigger pulses with a resolution of 12 ps for LEP beam instrumentation. The pulses produced are synchronous with the 352 MHz RF clock train, with a jitter of less than 5 ps, and are positioned with respect to the first RF ticks detected after the arrival of the 11.25 kHz turn clock pulse. The delays of the pulses are specified in two parts, first the number of RF ticks since the previous pulse (or first RF tick in the case of the first pulse in a turn), then an 8 bit fine delay (1 bit = 12 ps) that covers one RF clock period. The module is able to generate up to 32 pulses per LEP turn, but for the streak camera only 4 pulses separated by 21-23 μs are used, corresponding to the selection of 1 bunch out of every train.

These pulses are used to drive the fast sweep of the streak camera via the Timing Adaptor, as discussed in Section 3.1 above. As in the case of the slow sweep, the fast sweep has to be triggered in advance of the arrival of the photon bunch in the streak camera by a time that decreases from 60 to 8 ns as the sweep speed increases. The fact that the pulse for each fast sweep has its own fine delay with respect to the nearest RF clock makes it possible to individually adjust the position of the corresponding bunch image in the direction of the fast sweep.

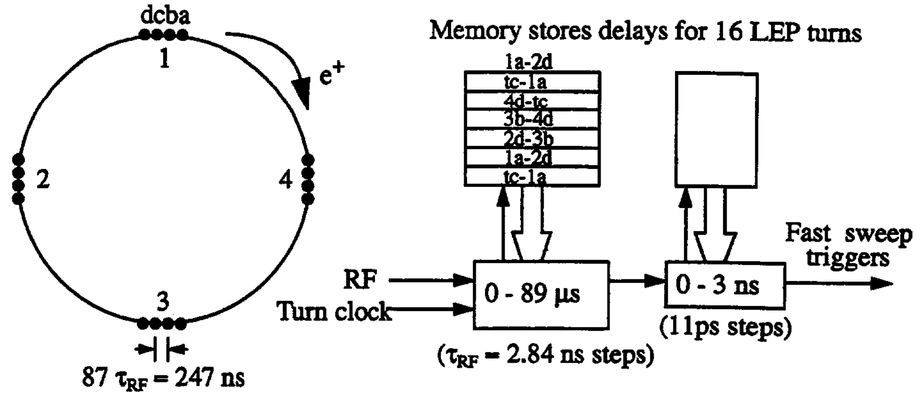

Figure 5 : Picosecond timing for bunch selection

When the streak camera fast sweep is triggered 4 times per LEP turn for a number of consecutive turns (i.e. a burst at 45 kHz), the bunch images of the first 16 or so turns appear displaced on the screen by up to 300 ps (see left part of Figure 6). This effect is reproducible and due to a small distortion of the fast sweep triggering in the first 1 - 2 ms of the burst. Rather than attempting to correct this distortion at the source, the LSD multi-page sequencer module has been programmed to present slightly different delays for the 4 trigger pulses in each of 16 consecutive LEP turns. Applying the same shifts for each trigger in a given turn produces the result shown on the right of Figure 6. An even finer correction has been implemented by having a variable shift for each of the 4 triggers per turn. This correction has only been tried with one fast sweep speed and it is quite possible that different corrections would be required for different sweep speeds. In the present implementation, the required delays are stored in the 16 memory pages of the LSD module and read each time the RF clock delay counter in the picosecond timing module passes through zero. It is also possible to disable the distortion correction by looping indefinitely on one memory page, which is the only way to generate more than 16 turns of triggers. The final version of the picosecond timing card, presently under development, will include similar functionality onboard but with 32 sets of 32 delays per turn.

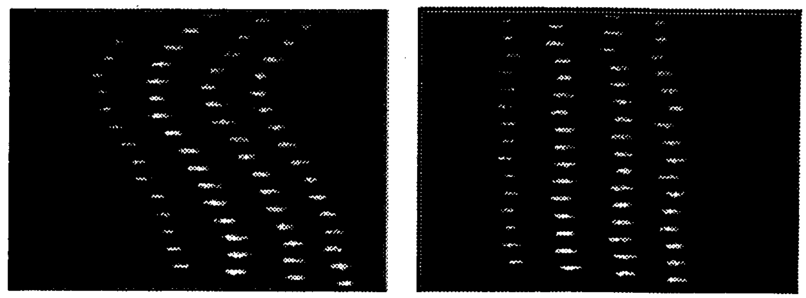

Figure 6 : Fast sweep trigger distortion before and after correction

3.3 BST mode

This alternate mode has been built into the present version of the Timing Adaptor, but has not yet been fully commissioned. The Beam Synchronous Timing system [11] is a dedicated message distribution system in which different beam instruments located at different points on the 27 km LEP ring can receive triggers within the same LEP turn. This means that exactly simultaneous measurements can be produced on different beam instruments. Different bits in the 56 bit message transmitted every LEP turn are allocated to different instruments; with some destined to trigger hardware directly, while others trigger interrupt driven software or contain data values. In the case of the streak camera, one "hardware" bit (which is either set or not during every LEP turn) is used to trigger sequences of measurements via the Timing Adaptor. It is possible to program the BST system to trigger any arbitrary sequence of turns, either continuously or linked to particular machine events, such as injection. For example, one could request images with 10 turns displayed with an interval of 8 turns between each displayed turn. Although the more exotic combinations of turns may not be very useful in practice, the possibility to vary the image frequency from 25Hz down to single shot is certainly an advantage of this mode.

In the BST Mode it is the arrival of the first BST streak camera bit, rather than the CCD ready, than enables the slow sweep trigger. The whole sequence of BST bits is then delayed 8 turns and serves to enable the appropriate sequence of optical gates and fast sweep triggers on the desired turns. The streak camera is then being driven by the BST, but asynchronously with the CCD camera, which would receive images from the camera phosphor screen at any moment in the CCD cycle. While this would mostly work for short bursts of turns (e.g. 11 turns = 1 ms), the results would not be reproducible, with some images lost and some with uneven light integration. To solve this problem, the manufacturer of the CCD camera was asked to provide a "CCD init" input that allows the restarting of the CCD frame clock at any moment. The first BST bit should then produce a "CCD init" pulse and then after an appropriate time start the sequence of streak camera triggers. Another problem with this mode of operation occurs in the Datacube image acquisition. When the CCD camera is free running, the phase of the video line frequency is maintained from one frame to the next, which allows the Datacube to use a PLL on the line synchronisation pulses to generate the pixel clock for digitization. However, the use of the "CCD init" breaks this phase (unless the "CCD init" is pulsed at an exact multiple of the line period) and the PLL becomes unusable. The Datacube permits synchronisation with an external pixel clock, but this should normally be continuous, i.e. even when there are no active pixels to digitize during blanking periods. However with some modification to the pixel clock of the CCD camera to add some extra pixel ticks between lines and appropriate programming of the Datacube input video parameters, it has been possible to drive the Datacube image acquisition reliably with the CCD pixel clock.

4. IMAGE PROCESSING

The MaxVideo20 pipeline image acquisition, processing and display card from Datacube forms the heart of the streak camera image processing system. Processing elements on the card can be reconfigured by software to form a data pipeline through which image data can then "flow". In the present implementation the MaxVideo20 card is only used to (i) digitize the CCD camera image in CCIR standard video mode with an 8 bit flash ADC and (ii) generate the final high resolution video display for the LEP control room. All the image analysis is done on a 68030 processor card running Microware's OS-9 operating system.

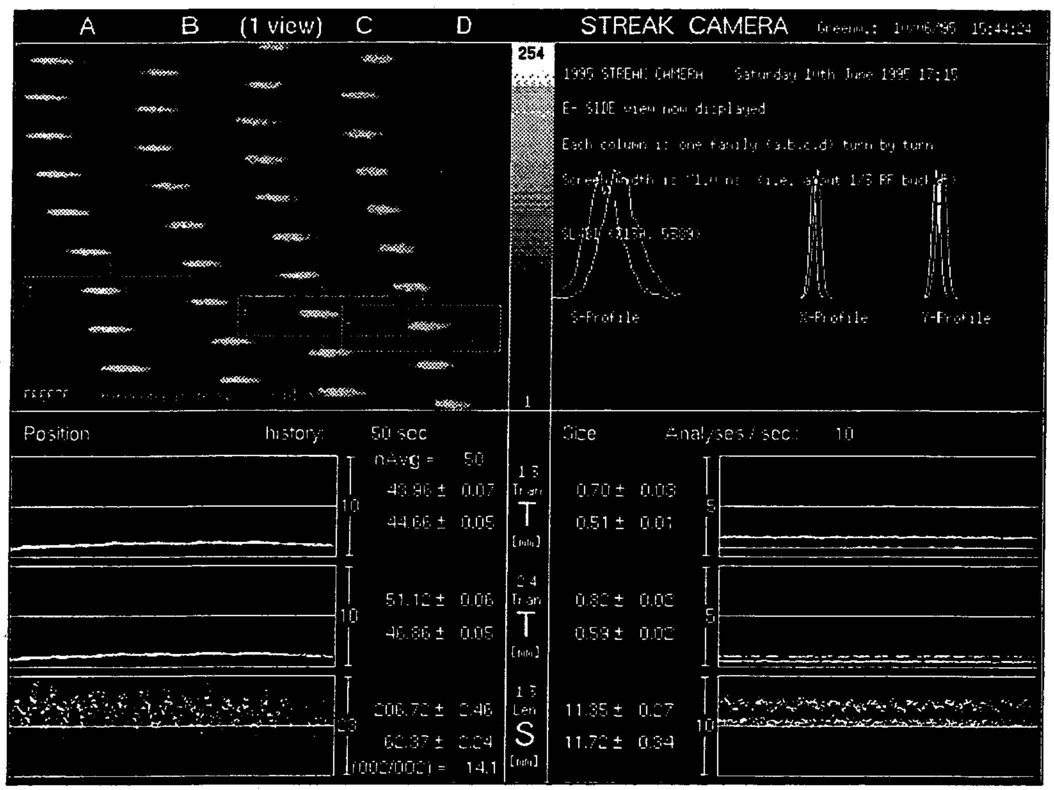

Figure 7 : High resolution video display transmitted to LEP control room

The image digitized from the CCD camera video output is refreshed at full speed (25 Hz in the case of normal "CCD mode" timing) in the top left part of the display. The intensity recorded in each CCD pixel is displayed in false colours, as indicated on the colour scale. The fast sweep appears to run from left to right and the slow sweep from top to bottom. In this example the side view of bunches 1a, 2b, 3c, and 4d of the e^- beam for 9 consecutive LEP turns has been selected. The width of the image corresponds to 1.3 ns, the height to about 0.8 ms. It can be seen that the distortion correction described in Section 3.2 has not yet been implemented.

The evaluation of bunch position and size is done on a copy of the image made by the MaxVideo20 card in 5 ms. This copy is only updated when the calculations on it are complete. When the screen copy indicated in Figure 7 was made, 4 bunches were being analysed at a frequency of 10 Hz. The rate is higher when the profiles shown in the top right part of the screen are not displayed. The 4 rectangles superimposed on the CCD image indicate the screen regions being used for the determination of the bunch positions and sizes. All rectangles are producing transverse measurements, while only rectangles 1 and 3 are measuring longitudinally. This was suitable for simultaneous measurements of the side and top views of the same bunch. For the measurement of bunch train families, all rectangles measure longitudinally and only 1 and 3 measure transversely. The bunch size calculation algorithm is quite complex but essentially consists of: (i) making projections of the bunch density distribution on the sides of the rectangle, (ii) estimating the background from the densities at the edges of the rectangle, and (iii) evaluating the bunch sizes from the standard deviations (σ) of the projections after background subtraction. The sides of each rectangle are at somewhat more than 4 σ from the centre of gravity of the enclosed bunch. The result of each evaluation of bunch position and size is converted to mm with conversion factors that depend on the sweep speeds and plotted on the corresponding continuously scrolling history graph. The time period shown in the history plot can be varied in steps of 10 s from 10 s to 24 h (50 s is selected in the figure). The individual values are also combined to produce an average and standard deviation for each measured quantity, and this is displayed numerically. The number of consecutive measurements used for this purpose can be varied from 17 to 3000 (50 is used in the figure). The average and standard deviation of each position and size is also available for reading out by a remote logging process (see Section 5.2). In the case of fewer than 4 columns of images on the screen it is possible to reduce the number of rectangles being used.

An important feature of the image processing is the continuous tracking of the bunch image by its enclosing rectangle. Each calculation of bunch position and size defines a new position and size for the rectangle. However in order that noise on the measurements does not generate oscillations in the rectangles, the calculated changes are strongly damped by limiting them to a one pixel change in position or size every analysis period. Nevertheless, this is sufficient for the rectangles to respond to real changes in the positions and sizes of the bunches. In the case of an abrupt change in the image, caused for example by a change in sweep speed, a given rectangle may lose the bunch it has been tracking. A complete screen bunch search algorithm enables it to locate the nearest one and lock onto it. This search may also be triggered at any time for all the rectangles by a software request to the system.

During acceleration in LEP the particle bunches change position in the RF bucket by a few hundred ps. This may result in some bunch images moving off the screen. An automatic procedure is available to keep the bunch in the first rectangle within a vertical band on the left side of the screen. When the bunch leaves this band in either direction a fine delay is changed that acts to bring it back. This fine delay is obtained by programming an 8 bit DAC on the MaxVideo20 card to add a DC level to the fast sweep trigger generated by the picosecond timing module, before sending it to the Timing Adaptor. The slope of the leading edge of the pulse is adjusted such that the pulse triggers about 1 ns earlier for full scale on the DAC. This automatic alignment procedure also compensates for any other effect that may change the bunch position on the screen, such as a slow drift in the 352 MHz RF clock distribution, or thermal effects in the streak camera or its associated electronics, as long as the effect is no larger than a few 100 ps.

The 800 x 600 pixel display produced by the MaxVideo20 at an interlaced frame frequency of 60 Hz requires a transmission bandwidth of 17 MHz, considerably in excess of the 5.5 - 7 MHz provided by the standard PAL transmission system. The separate analogue video signals for red, green, and blue (RGB), with synchronisation on the green channel, are therefore transmitted via 3 multimode optical fibre links over the 4 km to the LEP control room. Propagation delays on the 3 links are adjusted to be identical and industrial optoelectronic transmitters and receivers provide a bandwidth of 25 MHz.

5. Control Software

The streak camera system control architecture [12] is shown schematically in Figure 8, where the row of bubbles in the lower part of the figure represents the front-end control software resident in various VME computers, each controlling a part of the streak camera system described in the previous sections.

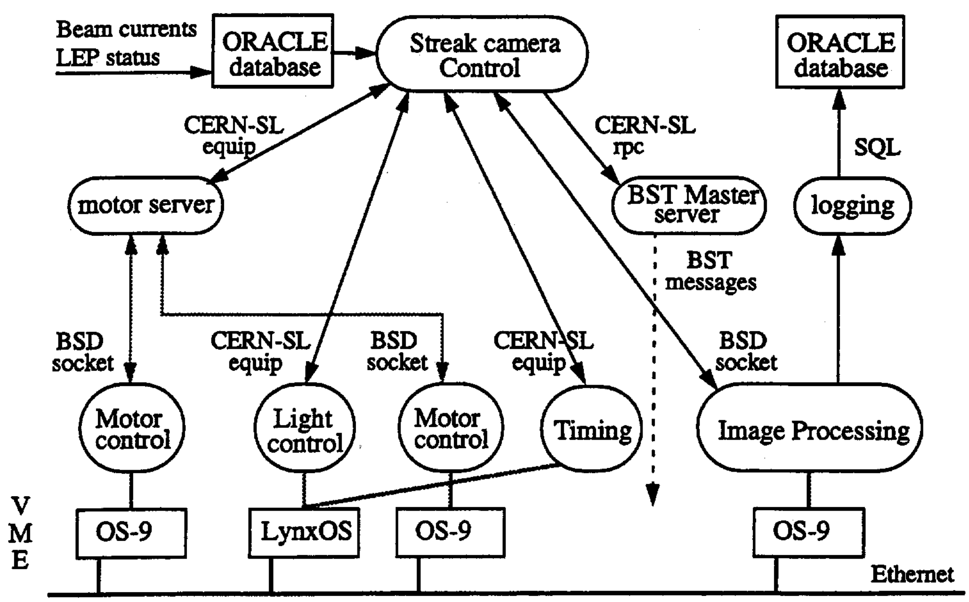

Figure 8 : Streak camera control architecture

These computers are diskless systems with 68030 processors directly connected to Ethernet, and their applications are loaded from a central file server via NFS (Network File System). The motor control software runs under Microware's OS-9 operating system and controls standard stepping motor interfaces. The light control and timing software runs in a single Lynx OS system, while the Datacube image processing system uses OS-9. For communication the "equip" package developed in the SL Division of CERN [13] has been used where possible. This is a layered data exchange protocol built on the CERN remote procedure call (rpc), which in turn uses a transport protocol handling datagrams. An "equip" server was developed for controlling the various VME modules of the LynxOS system, where it was installed as a thread sharing data with the acquisition thread that is used to run the light intensity control system. The motors were controlled via an "equip" server running on an HP-UX node that handles accesses to all such motor systems via TCP-IP using BSD sockets. TCP-IP is also used directly to communicate with the image processing system. In this case shared memory segments (MOPS [14]) visible to both the receiver process and the image analysis process are transferred, and a circular buffer is used to handle successive calls.

5.1 X-Window application

A Motif style graphical user interface (GUI) for control of the streak camera system is under development using the X-Window User Interface Management System (XUIMS) supported by the CERN SL Division Control Group [15]. This interface will replace the set of three presently used control applications, all of which have simple text interfaces. The main panel, as shown in the upper part of Figure 9, presents the current state of the view selection shutters, possible error conditions from the synchrotron light intensity control subsystem and relevant LEP machine parameters, in particular mini-wiggler magnet and beam currents which are retrieved from an on-line database. From this information one can determine at a glance whether there is a simple reason for the absence of images on the streak camera video display, e.g. no currents in the mini-wigglers used as the synchrotron light sources. Space has been reserved on this panel for a graphical representation of the evolution of the measured bunch lengths using the commercial plotting package XRT/Graph. This will allow monitoring of the correct functioning of the image analysis, even from places where the streak camera video display is not visible. If both top and side views are being measured, then the bunch lengths given will be corrected for the frontal spot size of the synchrotron light pulses [16].

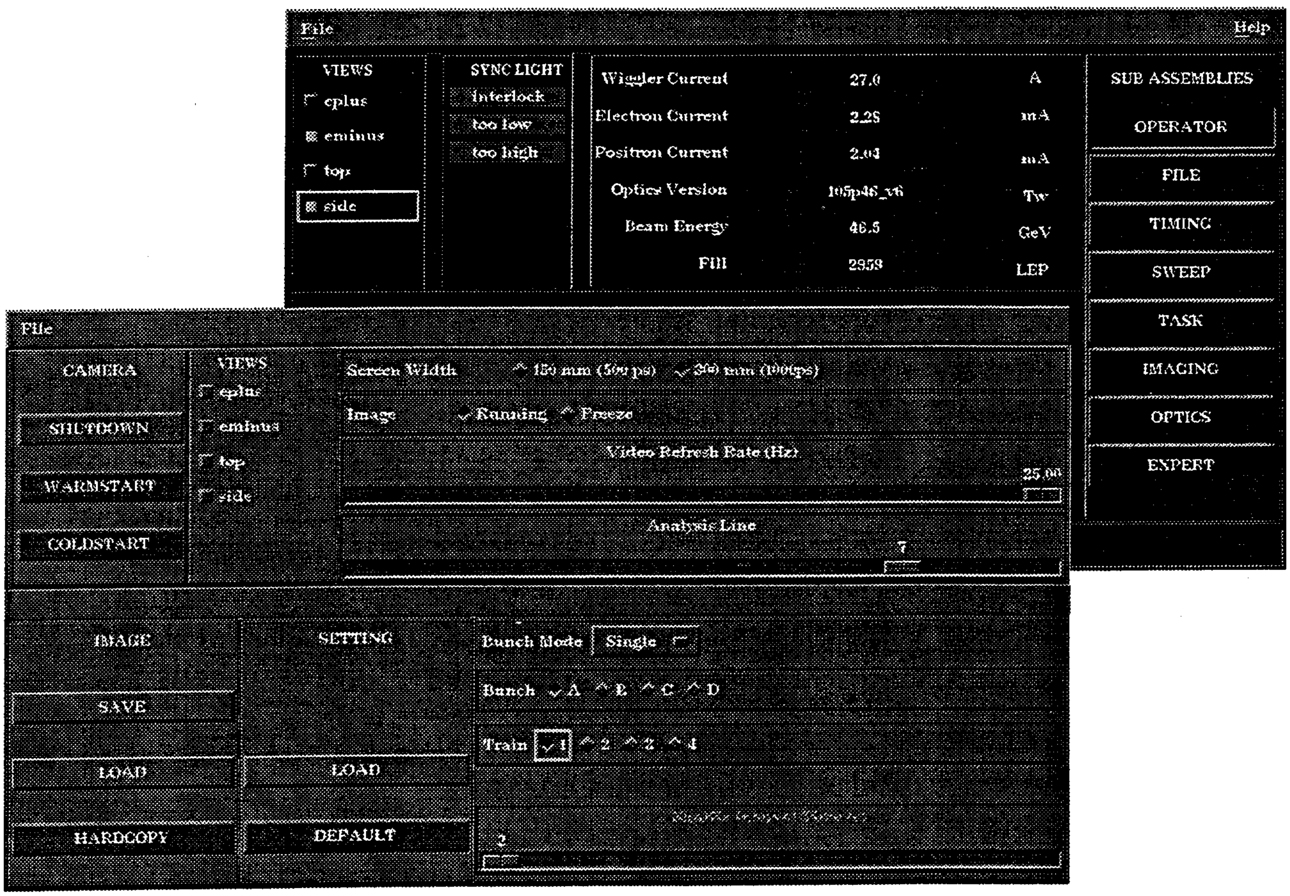

Figure 9 : Streak camera control main User Interface panel and sub-panel for LEP operators

Subsidiary panels group together all the controllable parameters according to type, e.g. timing, streak camera sweeps, motorised optical elements, etc. Control panels that are destined for use only by experts (e.g. adjustment of optical line elements) are password protected. A special panel is being prepared containing the limited set of functions considered to be necessary for the use of the streak camera by LEP machine operators, a first version of which is shown in the lower part of Figure 9. Starting in the top left corner of this panel, the first section is devoted to global operations on the system including "Shutdown", that ensures that no light reaches the streak camera; "Warmstart", that puts the camera once again into operation and "Coldstart", that allows the rebooting of the various low-level VME control computers. Then a "Views" section permits the opening and closing of the e^+, e^-, top and side shutters on the optical bench. The next section allows one to choose from the two most useful fast sweep settings, to freeze/unfreeze the CCD image display, to change the image refresh frequency (once the BST mode is commissioned) and to set the analysis rectangles on a particular displayed turn. In the bottom left section of the panel there are controls to save and restore the CCD image area of the video display to and from disk files and to make a colour or monochrome Postscript file or printed copy of the entire video display (such as that shown in Figure 7). Then there is a "Setting" section to allow the loading of predefined configurations (including the most useful beam, view, bunch and sweep combinations). Finally there a section for changing the bunch selections, but the operator would normally use the preset configurations proposed in the "Setting" section.

One important task of the top-level control application is to distribute relevant changes in the parameters of one sub-system to another. For example a change of either the slow or fast sweep (*motor control*) implies changes in trigger delays (*timing*), light attenuation (*light control*) and bunch dimension pixel to mm conversion factors (*image processing*). Changes to the identity and type of bunch images selected via the beam and view shutters and fast sweep timing must also be communicated to the image processing system (to keep the identification on the video display correct) and the logging process (to ensure that the correct data is updated in the database).

The introduction of the bunch train scheme and the interest in observing simultaneously the different bunch "families" (i.e. the different members of trains, labelled from "a" to "d") has considerably increased the number of beam/view/family combinations possible. Previously, only one bunch per turn was selected and both beams and views were displayed across the screen in the direction of the fast sweep. This produced 4 columns corresponding to e^+ top, e^+ side, e^- top, and e^- side. For bunch train operation in 1995 the default display consisted of 4 columns corresponding to one bunch family per train (usually 1a, 2b, 3c, 4d, as in Figure 7) corresponding to one beam and view. However combinations of e^+, e^-, top and side remain of interest, particularly with the likely use of only 2 bunches per train for LEP2 operation from November 1996, which reduces the number of beam/view/family combinations to 6 from the present 32 for 4 bunches per train.

5.2 Logging on measurement database

An Oracle database is used for storing the history of a large number of LEP beam and machine parameters. A logging process running on an HP-UX node reads the current particle beam, view and bunch selections and the results from the image analysis every 10 s. Embedded SQL calls then update the appropriate fields in tables (one for each beam) containing the last averages and standard deviations of position and size in all 3 dimensions for all possible 32 bunches (i.e. 384 values in all). This triggers another process which copies only the data which has been updated, tagged with a bunch identification, into the history tables. In this way only the measurements being generated by the streak camera system at any time (i.e. maximum 24 values, see Section 4) occupy space in the history tables, and data corresponding to a particular bunch or bunch family may be easily extracted for off-line analysis.

5.3 Configuration data handling

A system for handling configuration data stored in a single data file and required in different processes has been developed [17]. The identification of data items was inspired by the handling of resources in the X-Window System and consists of an ordered pair of alphanumeric tags. Routines are provided to retrieve, update and print the data. The data is stored in an ASCII text file and any comments added by a text editor are retained by the data update routine. As well as strings and decimal values, hexadecimal and boolean values can be included as required. Global variables of scalar type, structures and arrays can be transferred between the resource file and the program. In applications where data have to be stored in shared or dynamically allocated memory, data offsets are calculated during compilation and the absolute location of the start of the storage area is passed as an argument to the data transfer function at run time.

6. CONCLUSION

A streak camera is a rather specialised and delicate device which one is not normally required to control remotely. Nevertheless, a project was undertaken to make the LEP streak camera and its associated optics, timing and image processing available for use by operators in the LEP control room. A secondary goal was that the camera should provide regular measurements of bunch dimensions, in particular bunch length, for logging on a database. Several VME-based computers now provide all the necessary low-level functionality and the interface to the LEP control system. However, despite the successful adaptation of the timing, image processing and logging to handle the new bunch train scheme in LEP, serious hardware problems with the streak camera itself prevented the achievement of these goals in 1995. With the return to CERN of a renovated camera in spring 1996 and the completion of the X-Window control application, the 1996 LEP run should see the goal of a fully operational streak camera become a reality.

7. ACKNOWLEDGEMENTS

The project to convert the streak camera from the initial locally-operated experimental tool into a fully remote-controlled system usable from the LEP control room has covered a number of years. Although the authors of the present paper are now responsible for completing this project, many others have also contributed. Special mention must be made of G. Baribaud, who led the project in 1992-94 and contributed particularly to the light control, timing and video transmission [18] [19]. Two visitors, first Y. Solberg and then M. Werner, developed the image acquisition and processing software on the Datacube system. E. Rossa and a number of visitors have been responsible for the initial streak camera installation, optics and improvements to the original PC-based analysis software and have provided input throughout the project. E. Rossa has also supervised the design and implementation of the picosecond timing module by a visitor, P. Joudrier, and proposed its use for correcting the fast sweep trigger distortion. C. Beugnet provided the stepping motor hardware and G. Morpurgo developed the OS-9 motor control software, the HP-UX motor "equip" server and a library for programming the LSD module. P. Pivot contributed to the testing and development of the light control system. Finally, C. Bovet must be thanked for his continual support and encouragement of the project, as well as the modification of the synchrotron light source to accommodate the bunch train optics, in collaboration with M. Placidi.

8. REFERENCES

[1] C. Bovet, M. Placidi, "A dedicated synchrotron radiation source for LEP diagnostics", LEP note 532, April 1985.

[2] Manufactured by A.R.P, route de Hausbergen, 67309 Schiltigheim, France.

[3] E. Rossa et al., "Double sweep streak camera for LEP", Proceedings of the 2nd European Particle Accelerator Conference, Nice, June 1990, pp. 783-785.

[4] E. Rossa et al., "Real time measurement of bunch instabilities in LEP in three dimensions using a streak camera", Proceedings of the 3rd European Particle Accelerator Conference, Berlin, June 1992, pp. 144-146.

[5] R. Jung, "Precision emittance measurements in LEP with imaging telescopes, comparison with wire scanner and X-ray detector measurements", Proceedings of the International Workshop on Particle Dynamics in Accelerators "Emittance in Circular Accelerators", Tsukuba, Japan, November 1994.

[6] Datacube Inc., 4 Dearborn Road, Peabody, MA 01960, USA.

[7] C. Beugnet, "Description du module Stepping Motor Interface", CERN LEP/BI-MP/Note 89-8, November 1989.

[8] J.J. Savioz et al., "Le module Interlock Controller", documentation technique, CERN SL/Note 93-114 (BI), December 1993.

[9] J.J. Savioz, "Le module LEP Synchronous Delay", documentation technique, CERN SL/Note 95-03 (BI), January 1995.

[10] E. Rossa et al., "BE 3D", CERN SL/Note 91-01 (BI), February 1991.

[11] G. Baribaud et al., "The Beam Synchronous Timing Master", CERN/AT-CO/91-10 & CERN SL/91-28 (BI), July 1991.

[12] J. de Vries, "Streak camera software: present status and future strategy", CERN SL/Note 93-99 (BI), October 1993.

[13] P. Charrue, "Accessing equipment in the SPS and LEP controls infrastructure - The SL-EQUIP package", Software User Manual, CERN SL/Note 93-86 (CO), September 1993.

[14] W. Herr, "M.O.P.S. User Guide for C programs", CERN SPS/88-43 (AMS), January 1993.

[15] M. Vanden Eynden, "The X-Window User Interface Management System (XUIMS) at CERN", presented at this conference.

[16] F. Tecker, "Evaluation of the performance of a streak camera in beam diagnostics at LEP", diploma work PITHA 93/7, Physichalische Institute RWTH Aachen, Germany, November 1993, & CERN SL/Note 95-66 (BI), June 1995.

[17] J. de Vries, "Use of resource files for multiple clients sharing a single data file", in preparation.

[18] G. Baribaud et al., "Specifications for the operation of the first version of the streak camera from the main control room", CERN SL/Note 92-29 (BI), May 1992.

[19] G. Baribaud et al., "Remote control of the operational streak camera -The second phase", CERN SL/Note 93-80 (BI), August 1993.

FERMILAB-Conf-96/069

CHARGING CONTROL SYSTEM OF THE ENERGY CAPACITOR STORAGE OF THE LASER DEVICE ISKRA-5

A.N. Gascheev, I.V. Galachov, V.C. Bezuglov, V.M. Murugov
The Scientific Research Institute of Experimental Physics
Arsamas, Russian Federation.

I.V. Mozin, V.P. Gerasimov, A.M. Gusev, V.D. Fedorov
D.V. Efremov Scientific Research Institute of Electrophysical Apparatus
St. Petersburg 189631, Russian Federation

ABSTRACT

The ISKRA-5 pulsed laser device is designed for research in nuclear fusion. It has 12 channels working synchronously against a common load to support different experiments on thermonuclear target heating. For the main power supply it uses stored capacitor energy. The power used (67 MJ) is sufficient to cause fatal damage, so safety must be guaranteed. Therefore, the ensuring of power supply reliability is one of the main functions of the control system. This requirement can be satisfied by improving power supply viability, increasing the stability of capacitor bank charge and emergency energy discharge when one of the capacitor sections fails.

1. INTRODUCTION

The physical device ISKRA-5 was designed for the study of the interaction of high power laser radiation with matter in the fusion programme [1,2]. The main part of the device is an Iodine Laser in which a light pulse, formed by a master oscillator (MO) is multiplied and amplified in 12 similar channels. Each channel has 5 cascades of amplification. The total laser output energy is 30 kJ in a 0.25 ns.

The power supply for the source of light, producing excitation of the master oscillator and laser amplifiers, is a pulsed power system based on a capacitor bank. Thr total stored energy is 67.2 MJ and the operating voltage is 50 kV. For the combination of values of stored energy and pulse length, 10-35 µs, thr capacitor bank of ISKRA-5 has no equivalent in the world. Closest to it in parameters is that of the device NOVA (USA) which has a stored energy 60 MJ and pulse length of 600 µs. Such a level of stored energy can cause fatal damage to the power supply system and the device as a whole, so safety should be guaranteed. Ensuring the reliability of the power supply system is one of main functions of the control system.

2. FUNCTIONAL REQUIREMENTS

The structure of the ISKRA-5 laser and pulsed power system is shown in Fig.1. Each of the amplifier cascade A1-A4 has its own capacitor bank (CB-A1)-(CB-A4), to ensure that the independent operation of a cascade MO, with its capacitor bank CB-MO, is common to all channels. For laser work optimization, it is necessary that the capacitor banks of identical modules are charged to the same level of voltage, which can be changed. The charge system controls the energy storage and power line and contains some types of industrially available units. The system allows the control of charge and level of voltage on the storage.

The main charge parameters of various capacitor bank modules are shown in Table 1, the total number of such modules is 738, including those of the synchronization system.

Table 1

module type	operating voltage, kV	charge error, %	charge time, ms
amplifier	50	2	90
special generator	50	2	60
synchronization system	70	2	20

The power units are assembled as the usual circuit with step-up transformers and uncontrolled rectifiers. Feed control for the power units is provided by thyristor regulators.

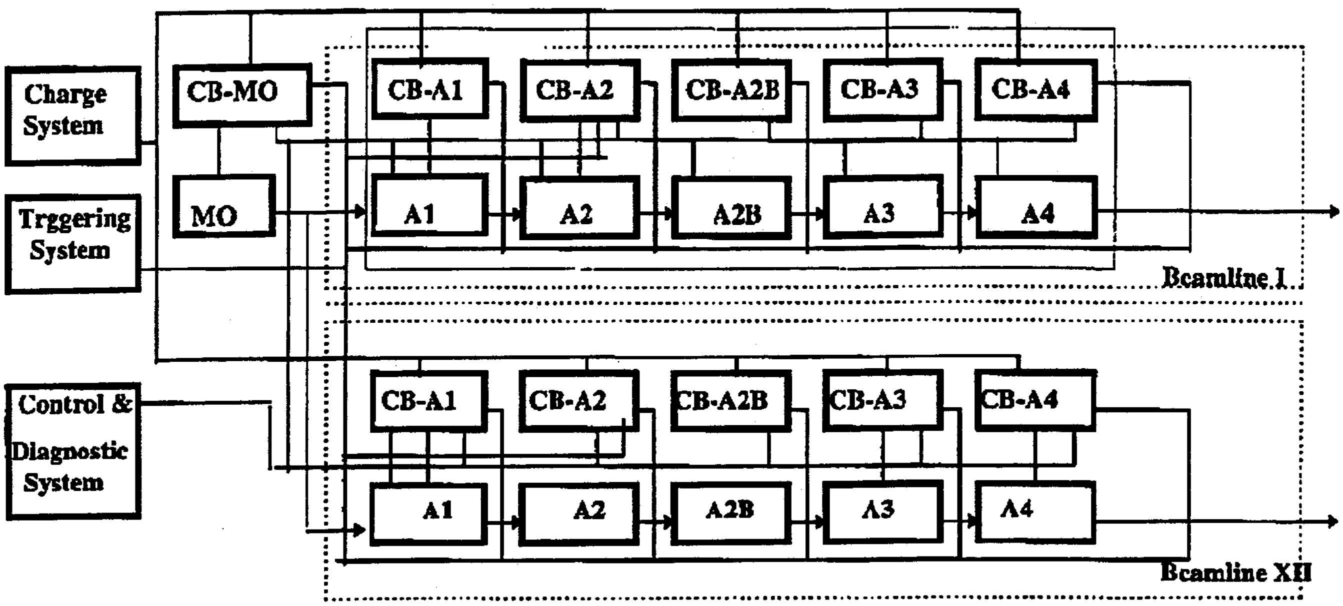

Fig.1. Structure of the ISKRA-5 laser and pulsed power system

The power supply circuit provides the following modes of capacitor bank charging:

- capacitor bank charging with fixed value of charge current,
- capacitor bank charging with fixed rate of voltage increase,
- combined charging mode.

When using the mode with a fixed value of current, the current is adjusted, as a rule, to the maximum value the power unit permits. The advantage of this mode is a stable load on the power unit, which is not changed even in the case of a short circuit. The shortcoming is that with a reduction of the bank capacitance, which is necessary for some experiments, the charging current is maintained. However all the elements of a capacitor bank module, to allow for the consequences of emergencies, should be designed on a complete charge current basis.

When charging the capacitor bank with a fixed rate of voltage increase, the charging current is changed depending on the capacitor load. This eliminates the shortcoming of the mode with fixed current. However, in case of an emergency short circuit, the charging current sharply increases and it is necessary to stop the operation.

As a rule, such a switching-off occurs as consequence of the operation disaster protection. Unfortunately we couldn't manage to obtain a good enough thyristor to maintain an output load with the levels of protection required.

Therefore a combined mode in the control of capacitor bank charging is preferable, in which a restriction of the sharp increase of charging current in an emergency situation is introduced simultaneously with a fixed rate of voltage increase. This avoids current overload and the necessity to stop charging. This mode is used with the incomplete composition of the capacitor bank. The control mode with a fixed value of charging current is chosen for experiments with the complete composition.

The control system should provide, together with conventional tasks:

- flexibility of storage charge control,
- remote and local control of storage charge using central and technological control panels,
- automatic charging,
- high-speed electronic protection,
- diagnostics of conditions of all technological systems, forecasting emergency conditions and the formation of control commands in extreme situations,
- localization and switching-off of the emergency part of the power supply while preserving the functioning of the device as a whole.

3. IMPLEMENTATION

The thyristor regulator of the input voltage for the power unit is designed using conventional circuit phase-pulse control. A high quality of auto-control is provided by a correction system. Some kinds of protection are provided, including high-speed protection against the maximum output voltage of the power unit and maximum charging current.

For elimination of the influence of emergencies in cells of a capacitor storage bank on the workings of the charging system, the mode with a fixed value of charging current is used. This permits the stabilization of a value of charging current with changes of load.

Particular attention was given to noise immunity of the control system. The general policy was to use electromagnetic relays not only as switches but also as logic elements. As it turned out, considerable distortion of the voltage waveform on the feeder cable occurs due to the large current drawn by the charging system. These distortions resulted in failure of the control system. Complete removal of this interference was obtained only by the transfer of the control system supply to another substation.

4. MAIN RESULTS OF OPERATION OF CHARGE CONTROL SYSTEM

Experience of operation with the energy storage has shown that it is possible to ensure reliable realization of experiments with a capacitor storage of 67.2 MJ. As was expected, in a number of experiments there were failures of some modules. Therefore the design of charging system provides a continuous capacitor bank charge despite failures occurring, which is stopped only by a command of the operator from a central control panel.

The technical decisions made in the design of the charging control system have allowed the realization of these aims. On failures of different kinds, including momentary short circuit of the power supply, there is no source switching-off and only a limiting of the charging current takes place. After switching-off a broken module with the help of a special high-speed safety fuse, the charging of the rest of the capacitor bank goes on in a given mode. The transients arising at failures are smoothed sufficiently and do not influence the charging process.

5. CONCLUSIONS

The operation of capacitor banks with a total stored energy of 67.2 MJ at an operating voltage of 50 kV allows us to make the following conclusions about the practicability of energy storage designs of similar and larger sizes.

1. With an increase of the value of storage energy a bank becomes more multielement and accordingly the probability of failure of some element is increased during bank charging or discharge. Thus, the failure of capacitor elements becomes a planned event, and appropriate measures for the reduction of such failures' influence on the efficiency of the bank and the device as a whole should be adopted.

2. The reliability of experiment realization is largely defined by the charge of the capacitor bank. Therefore it is very important to fulfill the following conditions during the charging stage:

- the tolarance of the charging to emergency operation,
- the localization of an emergency in any part of the charging system and its switching-off,
- the emergency discharge of the energy stored in the capacitor bank into a ballast load,

- the minimum time between the attainment of full voltage on the capacitor bank the issue of a firing command.

The design methodology for a high-energy capacitor bank and its control system which have been developed allows the creation of a capacitor bank of order of 100 MJ with the reliability required for experiments.

REFERENCES

1. G.A. Kirillov, V.M. Murugov, V.T. Punin, V.L. Shemyakin "High Power Laser System ISKRA-5" International Conference SPIE, OE/LASE'90, Los Angeles, CA, USA, January 1990.

2. V.G. Bezuglov, I.V. Galakhov, A.S. Gasheev et al. "The 67MJ, 50 kV Capacitor Bank for the High-Power Iodine Laser of the Facility ISKRA-5: Description and Experience" 10 th IEEE International Pulsed Power Conference, Albuquerque, USA, July 1995.

The New Vacuum Control System of the CERN PS Complex

R. Gavaggio, P.M. Strubin CERN / AT-VA

Between 1992 and 1995, a new vacuum control system has been installed step-by-step on the different accelerators of the CERN PS Complex (Linac2, Linac3, PS Booster and PS) as well as on the related beam transfer lines. This project has been carried out in parallel with the installation of the new main control system of the CERN PS accelerators.

The initial goal was, on one hand, to rejuvenate old equipment whose technology was very difficult to adapt to the new control methods, and on the other hand, for obvious compatibility and maintenance reasons, to standardise the hardware and software of these different systems. The hardware's two level architecture fits in very well with our needs of distributed control. Standardised and industrial equipment (VME and G64 boards and chassis, RS232 and X25 field buses, etc.) has been used where ever possible. The software is structured in two layers to obtain a standard interface to all the vacuum equipment for each of the different accelerators. This presentation describes the hardware and software architecture of this project as well as the different equipment used.

1 INTRODUCTION

The CERN PS complex is a set of inter-connected accelerators which can accelerate different types of particles and send them on to other accelerators (AAC, LEAR, SPS and LEP) or directly to experimental areas. It consists of the following accelerators: LINAC2 (Proton Linac), LINAC3 (Heavy Ion Linac), PS (Proton Synchroton), PSB (PS Booster), LIL (Lep Injector Linac), EPA (Electron Positron Accumulator) and all the transfer lines between these different accelerators.

In 1992 the controls group of the PS division started the important project of rejuvenating, step by step, the main control system of these different accelerators [1]. At the same time the vacuum group of the AT division took a similar action to install its new control system. The main goal was to rejuvenate old equipment which did not adapt well to the new control methods, and also to standardise the hardware and software of all these accelerators using standardised and industrial equipment where ever possible.

The new vacuum control system became operational on the accelerators as follows: in 1993 LINAC2, in 1994 LINAC3 and PSB, in 1995 PS and the transfer lines. At the current time, this represents about 700 pieces of vacuum equipment (pumps, gauges, valves, etc.) accessible via the new control system.

This paper gives an overview of the new vacuum control system and describes the different equipment and software architecture used.

2 HARDWARE LAYOUT

The vacuum level on the different accelerators and transfer lines ranges from 10^{-6} to 10^{-12} mbar. This requires a large variety of measuring and pumping equipment, the main types being :

- Gauges (Pirani, Penning, etc.) ≈ 180 units.
- Ion pumps (from 50l/s to 1000l/s) ≈ 300 units.
- Sublimation pumps ≈ 130 units.
- Sector valves ≈ 60 units.
- Gas analysers ≈ 10 units.
- Pumping stations ≈ 80 units.

At the present time each set or piece of equipment on the control system has its own chassis for local control plus an RS232 port for remote control. The pumping stations only run in stand-alone mode. The RS232 interface has been selected because it is a widely used standard which allows us to easily integrate equipment from different sources (CERN, industrial manufacturers, etc).

The general layout of the system is exactly the same for each machine (fig. 1). It is an architecture with two layers which are connected together by a local network based on X25/RS232. The upper layer is composed of a VME crate (called DSC) used for data acquisition and communication with the main control system of the accelerator. The lower layer consists of all the G64 control chassis and industrial equipment.

2.1 VME Chassis (DSC)

This chassis, provided and maintained by the PS-CO group, is used for the acquisition of all the vacuum data (pressures, status, etc.). The connection with the upper level of the main control system of the accelerator is made via an Ethernet network. The DSC consists of a 6 slot, diskless VME chassis with at least the following cards:

- The CPU unit, type Motorola MVME147, with eight MB of memory and an Ethernet interface.
- The SAC unit used for remote reset and supervision.
- The X25 unit(s), type Motorola MVME336, with six 1 MB/s full duplex data link ports.

2.2 X25/RS232 Local Network

The X25 local network is used to multiplex data from many RS232 channels. It transfers data over distances up to 300 meters. One VME module allows us to connect, via a hub module, up to six remote terminal servers providing 16 full duplex RS232 lines each, that is to say 96 lines per VME card.

2.3 G64 Chassis

The old vacuum control system used many specific chassis, different from one machine to another, and from different technologies. The rejuvenation of the system has allowed us to limit the number of specific chassis to only 3 different types. These chassis, developed at CERN, use a 3U Europa frame and a G64 bus partially modified to get a direct access, via the bus connector to the inputs/outputs. One chassis accepts up to 4 standard G64 cards and up to 8 specific cards (160 or 220 mm length).

Each chassis has a CPU card (2Mhz MC6809E) with 32 kB of RAM and EPROM memory, a Real Time Clock circuit and two RS232 interfaces. The first one is used for communicating with the DSC, while the second can be used locally with a terminal to execute tests and local control.

The 3 types of chassis are the following:

- The Valves Control Unit (VCU) provides local and remote control for up to 8 valves of any type used at CERN. It can also be used to drive industrial valve controllers. In addition it handles all the safety interlocks used to close the valves in case of abnormal conditions (leaks, power-off, etc.).

- The Pumps Control Unit (PCU) is used for the remote control, status and pressure acquisition of up to 8 Ion pump power supplies of any type. In particular it allows us to switch on the selected pumps in sequence to avoid overloading.

- The Sublimation Control Unit (SCU) drives up to 16 Sublimation pump power supplies. It can store in local memory all the pump parameters (e.g. number of sublimations on each filament). It can also start and handle sublimation cycles for each pump individually. The same chassis can be also used to control 8 Ion gauge power supplies.

2.4 Industrial Equipment

At the present time the trend is to use more and more industrial equipment from the vacuum industry and to decrease the number of specific chassis developed at CERN.

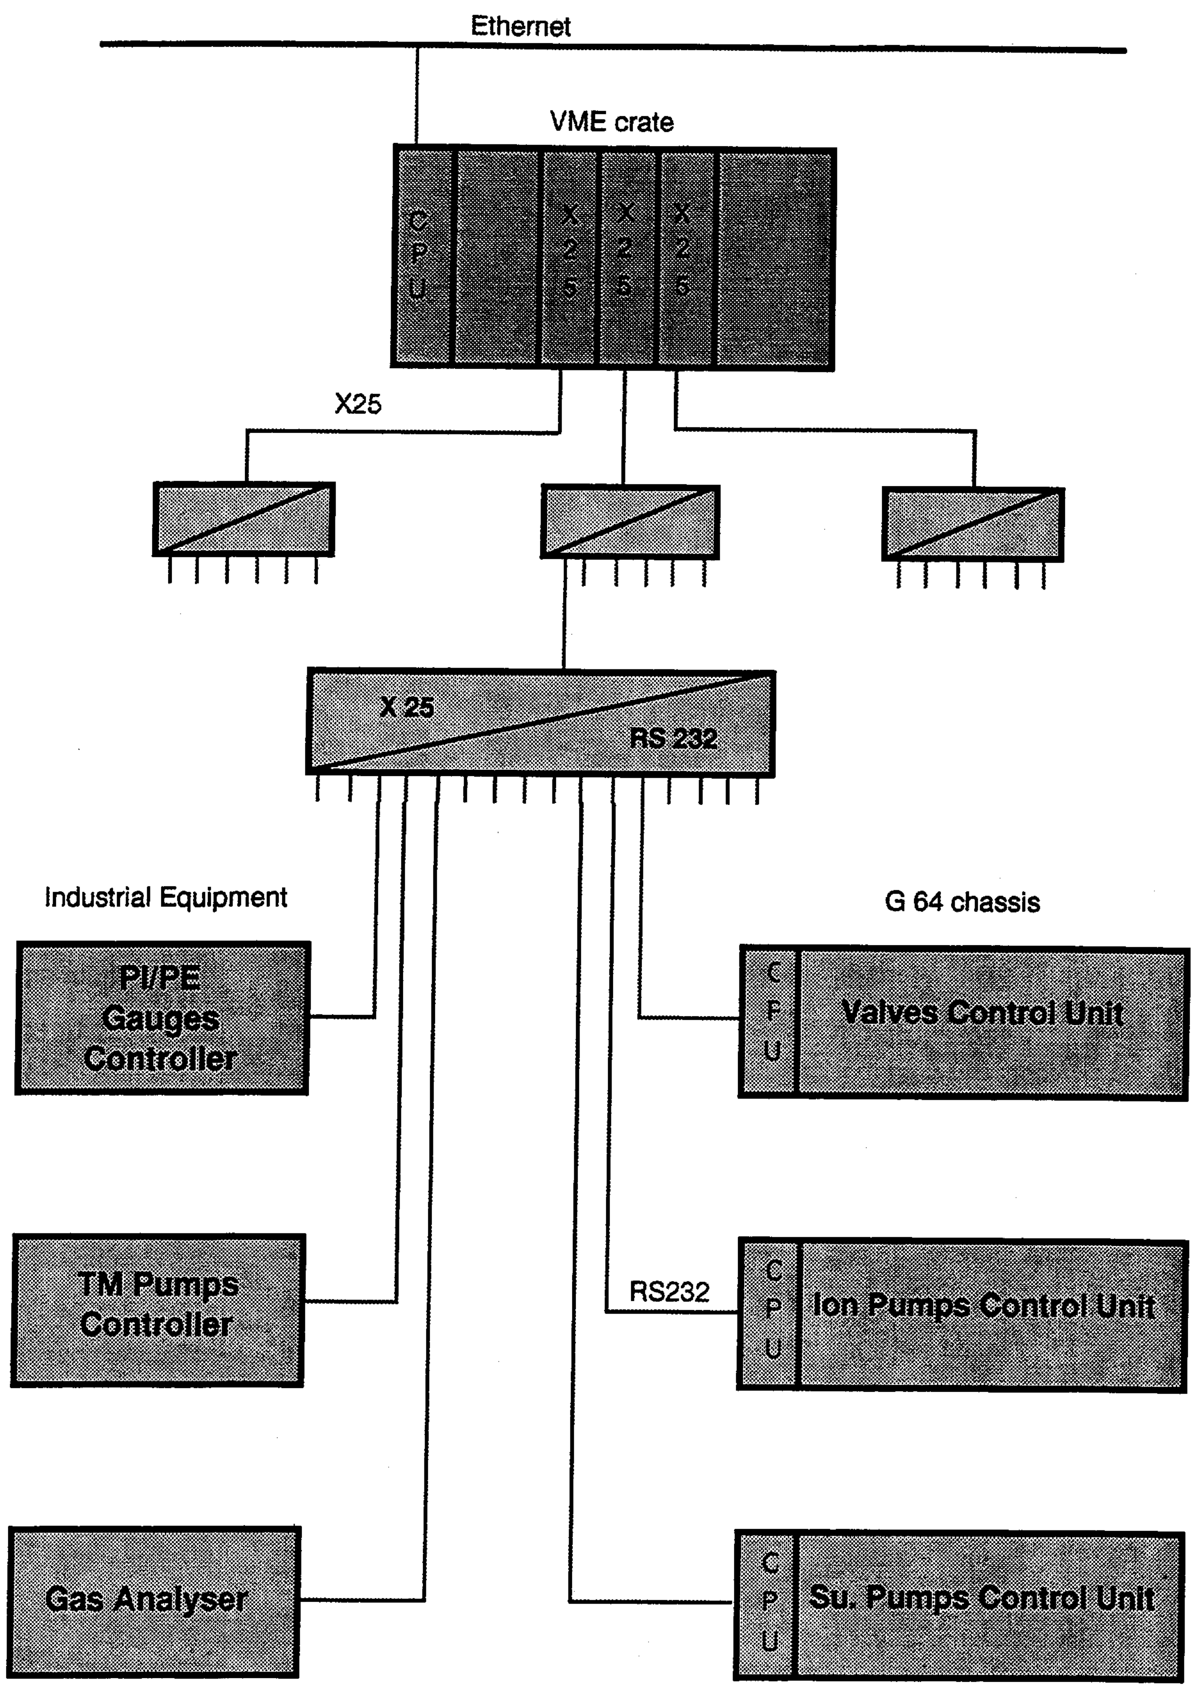

Fig. 1. Overall Layout

The most frequently used industrial components are:

- Pirani , Penning, Bayard-Alpert, etc. gauge controllers.
- Turbo-molecular and ion pump power supplies.
- Gas analysers.

Each piece of equipment has, as standard or as an option, an RS232 interface which allows easy integration into our control system.

3 SOFTWARE ARCHITECTURE

The main control system for the accelerators in the CERN PS complex uses a 3 layer software architecture:

- The upper layer, running at the work-stations level, includes all the application programmes [2], the logging, the archiving of the data and the databases.

- The middle layer includes all the programmes in the DSC. This layer is split into two parts :

    * Part 1: the interface with the upper layer (servers, equipment modules, etc.).
    * Part 2: the specific software, the interface and the drivers which communicate with the lower layer.

- The lower layer running in the specific and industrial equipment.

The software for part 1 and the upper layer is the responsibility of the PS-CO group, while the rest of the software, part 2 and below, is the responsibility of the user groups (e.g. vacuum).

Therefore, at the vacuum system level, we find a two layer structure (DSC level and G64 level) which uses all the local processing properties (distributed intelligence) of our equipment.The software uses an operational protocol for the vacuum systems based on the use of standard models of our equipment (one model per equipment family: gauges, valves, pumps, etc.) [3], [4].

3.1 Specific Software in the DSC

The programme running under the operating system Lynx OS has the following functions:

- Interface with the main control system.
- Message handling.
- Communication with vacuum equipment.

3.1.1 Interface with the Main Control System

The interface with the main control system is made through a queue mechanism (Sys V queues) in which a CERN standard protocol called USAP, Uniformisation of Software Access Procedures [5], is used for message passing from the equipment modules to the specific software. The acquisition and control messages can be passed synchronously or asynchronously.

3.1.2 Message Handling

The received messages are checked, compacted and formatted to be sent to the specified equipment. There are two different cases:

- The equipment is controlled by a G64 chassis. In this case the message is sent directly.
- The equipment is of an industrial type. In this case the message is adapted to support the specific protocol of the manufacturer.

The answers to these messages are handled in the same way before they are sent to the upper layer.The use of a local database, downloaded at the boot of the DSC and which stores all the parameters of the equipment, allows us to have entirely data-driven programmes.

3.1.3 Communication with Vacuum Equipment

A programme is used for the transmission of all these messages via a unique X25 driver. It handles the synchronisation, error detection, time-out handling, etc.

3.2 Software in the G64 Chassis

The software structure is exactly the same, whatever the type of G64 chassis, and uses a completely data-driven concept.

The embedded program, written in Pascal and Assembler, is composed of 4 main parts :

- The Real Time part which controls the vacuum devices (actuations, data and status acquisitions). It is the only specific part of the program.
- The Data Tables which contain all the information about the vacuum devices (number, type, parameters, etc.).
- The Communication Handler which receives and sends back messages from and to the DSC. The structure of the messages is compatible with the USAP frame but is adapted to our specific needs.
- The Local Control and Test program which is a set of facilities, easy to use by the specialist, for a direct access to the equipment.

4 CONCLUSION

After three years of use, this new vacuum control system gives full satisfaction. Its modularity, hardware and software, allows us to cope easily with modifications to or replacement of equipment. Furthermore, the limited number of types of chassis and their compatibility on all the CERN PS accelerators, make the maintenance easier.

The next and last step will be the installation, in the beginning of 1996, of this new vacuum control system on the LIL and EPA accelerators. This will raise the number of pieces of vacuum equipment accessible through the system to about 1000.

ACKNOWLEDGEMENTS

We would like to thank R. Ould-Saada and M. Steffensen who have developed the software at the DSC level, J-P. Bertuzzi, B. Moine, F. Morcel and E. Siebert who have installed and commissioned all the equipment.

We want to include in these acknowledgements our colleagues of the PS Controls Group for their effective collaboration.

REFERENCES

[1] F. Perriollat and C. Serre, The new CERN PS control system overview and status, ICALEPS'93 Berlin.

[2] J. Cuperus, F. Di Maio, C.H. Sicard, The operator interface to the equipment of the CERN PS accelerators, ICALEPCS'93 Berlin.

[3] P. Strubin et all, Operational protocol for vacuum system, CERN / AT-VA 91-6.

[4] R. Gavaggio, M. Steffensen, P. Strubin, Using models to allow for object oriented programming of the vacuum control systems, EPAC'94 London.

[5] G. Benincasa et all, Final report on the uniform equipment access at CERN, CERN / PS 93-16.

Controls Upgrade of the RF Systems of the CERN PS

R. Garoby, M. Gourber-Pace, S. Hancock, F. di Maio, N. de Metz-Noblat
PS Division, CERN, CH-1211, Geneva 23, Switzerland

Abstract

The latest upgrade of the control systems of the CERN PS Complex has seen the replacement of all components of the controls interface to the RF systems of the PS machine. New hardware offers features which have been exploited not only by concomitant software developments but also by changes in the philosophy underlying the operation of those RF systems. In addition, a novel measurement scheme has been installed which provides a display of RF parameters by processing a wide variety of signals continuously sampled at 1 kHz. New timing diagnostics have also been implemented.

This paper reviews the main advantages of the new equipment and the evolution in operational principles which it has afforded.

1. Introduction

The PS (Proton Synchrotron) Complex comprises three linear accelerators and six circular ones, all interlinked. At its heart lies the thirty-six year old PS machine which handles a variety of particle beams on a cycle-to-cycle (so-called "PPM") basis. A rejuvenation of the control system of the PS is foreseen in two stages, the first of which has been completed and has affected all components of the interface to the three RF systems of the machine.

Figure 1 is a greatly simplified schematic diagram showing the main features of a PS RF system. Three types of control variable are employed: analogue functions (RF voltage, servo loop gains, etc.); timing pulses; and bit patterns for quasi-static parameters which are only refreshed once per machine cycle. The corresponding hardware comprises a total of 6 VME, 2 CAMAC and 21 G64 crates.

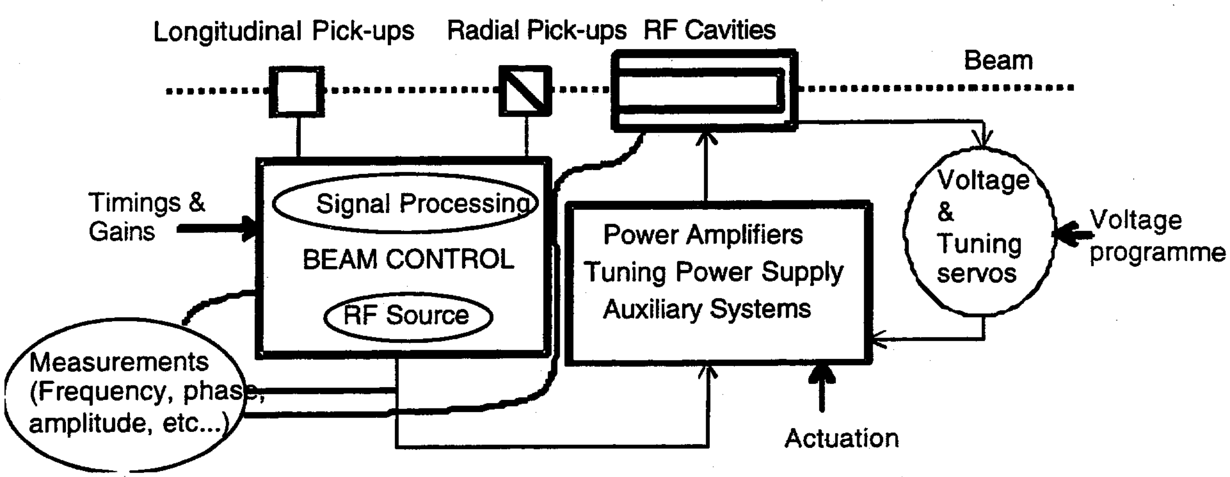

Figure 1. Basic RF layout.

2. Hardware

2.1 Function Generators

New VME-based function generators [Ref. 1] have increased by an order of magnitude the number of vectors available for the production of analogue signals. The need for extensive combinations of generators and summing amplifiers to provide detailed functions has thus been removed. However, the tremendous flexibility of the principal RF system of the PS derives from a combinatorial philosophy which has, therefore, been extended in this particular case. In order to avoid the transmission of analogue signals over long distances, a galvanically-isolated serial link carries the instantaneous digital values of the generated function either to a simple DAC for normal applications or to a multiplier plus DAC which converts the product of two such digital inputs. The latter permits different cavity voltage programmes to be generated within the same machine cycle as variants of a common fundamental programme.

In addition to more vectors, a function may contain multiple internal stops to hold the output at a certain value until a restart pulse is received and processing of the vector table is resumed. This facility is used extensively to synchronise transitions in the individual cavity voltage programmes with each other and with specific events (e.g., the opening or closing of the corresponding cavity short-circuits). The overall number of RF timings has been reduced since the wealth of vectors allows the duration of a flat first vector to determine the effective start of many functions. This trades off a dedicated start pulse against just one extra vector. Further economies have been possible because PPM means that restart pulses may be "re-used" at different times on different machine cycles.

A function may also be regenerated *ad infinitum* by processing the vector table in a loop. This permits the generation of a low frequency (< 20 kHz as the minimum vector duration is 5 μs) periodic function burst for modulation purposes from a description of one or more periods in the vector table.

2.2 *Timing Generators*

New VME-based hardware has permitted a change in the philosophy of timing generation. The majority of RF timing pulses are now generated in absolute time (with respect to the start of a machine cycle) rather than as delays from other timings. This is an advantage in large timing systems since there are no cabled links. Indeed, the very notion of timing structure has merely been emulated by establishing logical links between timing channels. A hierarchical timing cascade, which can be structurally different from cycle to cycle, may be maintained at application program level even though the hardware generates pulses entirely independent of each other. A further advantage is that a negative interval may be programmed between a linked timing and its logical predecessor without the need for an advanced prepulse (see Figure 2).

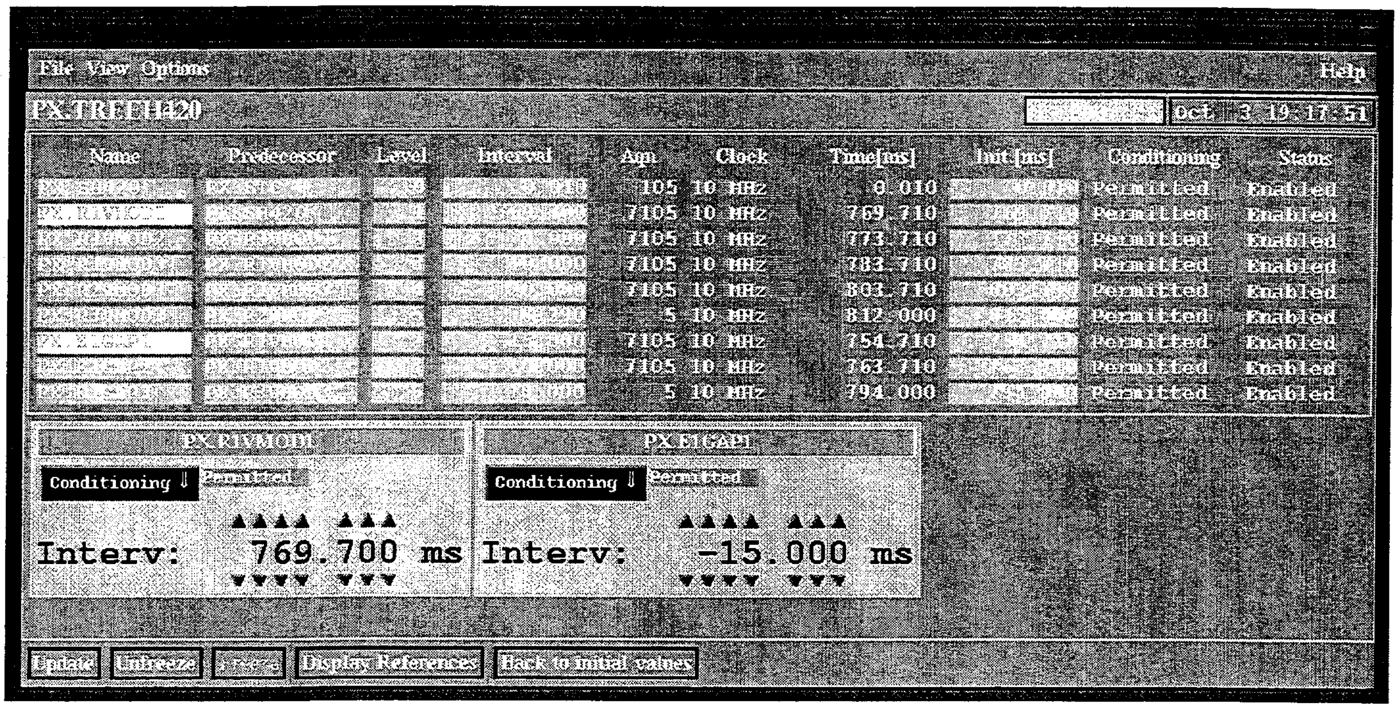

Figure 2. Linked timing application program.

Underlying the new system [Ref. 2] is a master timing generator (MTG) which distrubutes timing information on a single cable to all timing receiver (Tg8) modules [Ref. 3]. Each Tg8 receives the information in the form of timing frames which may subsequently initiate locally programmed actions. An action may add a delay (specified as a count of either an internal [10 MHz] or external clock train) before producing an output pulse, a VME interrupt, or both. Different actions may pertain to the same output channel and this feature is exploited to produce multiple restart pulses for function generators. Up to eight timing frames are transmitted during a 1 ms time slot and are validated by a "1 kHz event" at the end of this time. Thus, absolute time is readily described in terms of a number of Tg8 internal clock pulses after a particular millisecond from the start of a machine cycle. Time delays may be expressed with respect either to some other MTG event or to an external start.

The 1 kHz frequency of the MTG is derived from a 10 MHz rubidium gas oscillator which is, in turn, updated by a radio link from an atomic clock. However, the encoding of frames by the MTG and their decoding by a Tg8 are performed using free-running 4 MHz oscillators with the result that the 1 kHz events are subject to a jitter of ±250 ns. This is entirely sufficient for all except a small number of beam-related timings which require external starts and external clock trains.

2.3 High-power Equipment Interface

The high-power drive electronics of the twenty-one cavities in the PS are extensively protected by interlocks. The actuation of these systems is controlled and their status monitored via a completely new, G64-based interface [Ref. 4].

3. DIAGNOSTICS

3.1 Sampled RF Measurements

A powerful digital system has been assembled which provides a measurement of the essential RF parameters each millisecond during the active part of every machine cycle. The 1 kHz sampling frequency is derived from the same oscillator on which the MTG is synchronized. Various real-time tasks control the acquisition of: the beam revolution frequency and selected RF frequencies; the phase with respect to a revolution frequency reference of each of the principal RF cavities; the harmonic number of each of those cavities and all programmed and detected cavity voltages. These data are made available during the dead time at the end of a cycle for processing by an additional real-time task [Ref. 5] which computes the resultant voltage and phase components of the principal RF system by summing vectorially over all the cavities which are on the same harmonic (see Figure 3). It also computes the phase error, with respect to the phase sum for the appropriate harmonic, of each principal cavity. Both the raw and processed data can be accessed for display or further treatment.

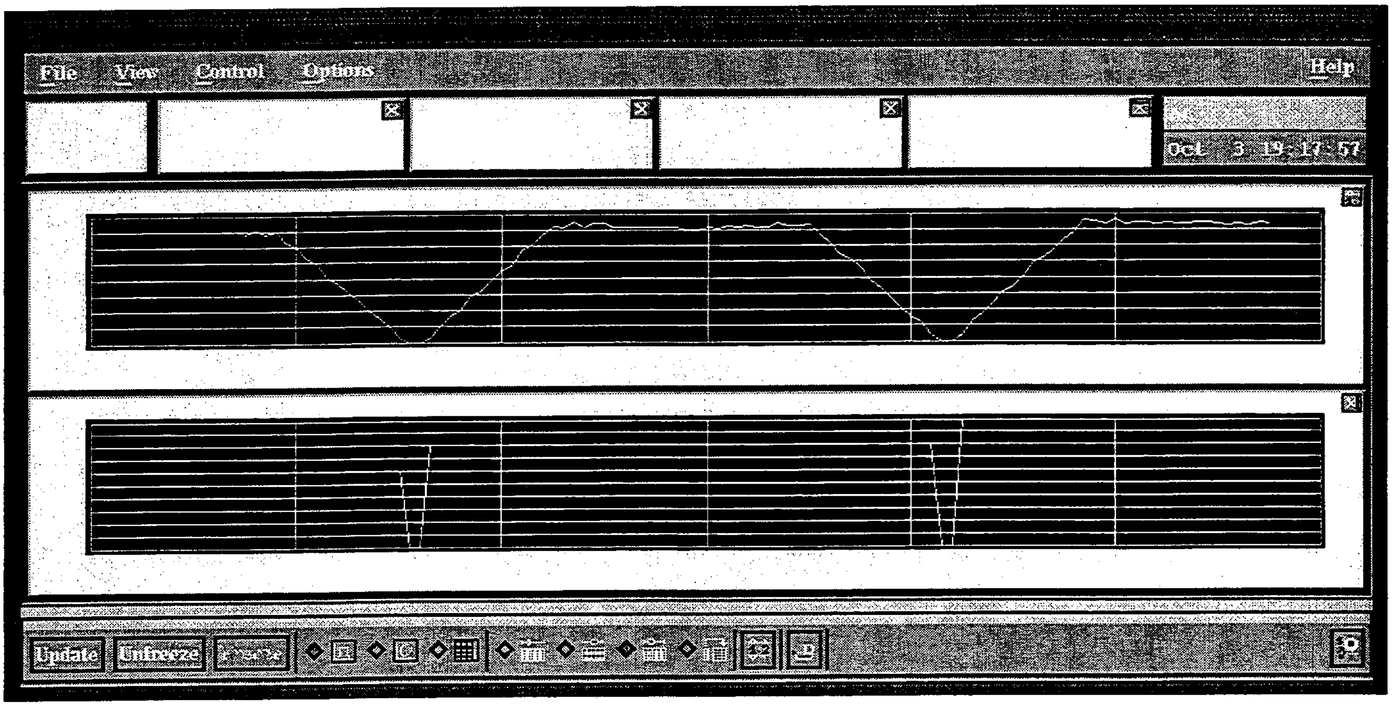

Figure 3. Total detected voltage and corresponding harmonic number measurements.

3.2 Timing Surveillance

Among the most difficult tasks in the operation of an accelerator is the diagnosis of timing problems. A timing surveillance (TSM) system [Ref. 6] has been developed to monitor simultaneously some 150 timing channels. All pulses, including multiple ones, are recorded with a resolution of 100 ns and an exhaustive list may be produced for an entire supercycle of the machine.

The precision of the TSM measurement is a consequence of its reliance on the widely-distributed 10 MHz clock described above. However, it is insufficient to monitor the beam-related timings that are required to synchronize the PS with its neighbours and client machines for injection and extraction purposes. To this end, a transfer timing surveillance (TTSM) system [Ref. 7] has been developed to measure time differences between key signals with a resolution of 1 ns. Eight different types of beam transfer may be monitored and warnings are issued for tiings that are out of tolerance.

3.3 Fault History

Any problems with the high-power RF systems of the machine are automatically logged so that a complete record of faults [Ref. 8] is available to specialists.

4. ACKNOWLEDGEMENTS

The success of the controls upgrade of the RF systems of the PS machine is due in large part to the dedicated technical support of P. Maesen, J.-P. Terrier and J.-L. Vallet. Important software contributions were made by A. Campbell, J.-M. Nonglaton and A. Gagnaire.

5. REFERENCES

[1] W. Heinze, R. Maccaferri, "GFAS: a new simple VME function generator", CERN/PS/CO Internal Note 93-09 (Tech.) (1993).

[2] J. Lewis, V. Sikolenko, "The new CERN PS timing system", ICALEPCS, Berlin, Germany, Oct. 18-23, 1993, Nucl. Instr. and Meth. A352 (1994), p. 91.

[3] G. Beetham, G. Daems, J. Lewis, B. Puccio, "A new VME timing module: Tg8", in Proc. ICALEPCS '91, KEK Proc. 92-15, KEK, Tsukuba (1992), p. 360.

[4] P. Maesen, "The G64 interface used to control the radio-frequency equipment in the PS", CERN/PS/RF/Note 95-12 (1995).

[5] A. Campbell, "Real-time formulae processing task for the PS 10 MHz cavity measurement system", CERN/PS/RF/Note 95-28 (1995).

[6] P. Bobbio, "Timing surveillance module", CERN/PS/CO/Note 91-07 (1991).

[7] S. Payaud, "Transfer timing survey module", CERN/PS/RF/Note 93-20 (1993).

[8] M. Beharrell, G. Benincasa, J.-M. Bouché, J. Cuperus, M. Lelaizant, L. Mérard, "Model-based, detailed fault analysis in the CERN PS Complex equipment", ICALEPCS'95 in these Proceedings.

Experimental Results on the Design for the APS PID Global Orbit Control System

Jeffrey A. Kirchman and Youngsoo Chung
Argonne National Laboratory
James P. Bobis
Northern Illinois University

Abstract The Advanced Photon Source third generation synchrotron light source needs a stabilized particle beam position to produce high brightness and low emittance radiation. Global orbit correction control is introduced and is utilized to satisfy the demanding needs of the accelerator. This paper presents the experimental results for determining an effective and optimal controller to meet the global orbit correction requirements. These requirements include frequency/time domain demands consisting of vibrational noise attenuation, limiting of controller gains for stability and improving the system time response. Experiments were conducted with a digital signal processor implementing various PID sets to make comparisons between simulations and experiments. Measurements at these PID sets supported the results of software simulation.

I INTRODUCTION

The 7-GeV Advanced Photon Source (APS) at Argonne National Laboratory (ANL) utilizes state of the art technology to produce the brightest beam of high-energy x-rays available for research purposes. The APS storage ring requires stabilization of the particle beam position to achieve the low emittance and high brightness radiation critical to a third generation light source. To achieve this, global and local beam orbit correction feedback is employed [1]. The global feedback control system consists of a PID controller in a digital communication system that collects beam position monitor (BPM) data in digital form, transmits it to a digital signal processor (DSP) and then forwards the resulting power supply correction currents calculated by the DSP to the relevant correctors. This paper presents the experiments that were conducted with a DSP implementing various PID sets to make comparisons between simulations [2] and experiments.

II REQUIREMENTS

The block diagram of the APS closed loop feedback system is shown in Figure 1. The low pass filter (LPF) is designed to provide system stability and to reduce feedback from the noise source. The controller is the PID. The compensation filter (CF) is a derived transfer function that attempts to neutralize the eddy current effects of the magnet and vacuum chamber.

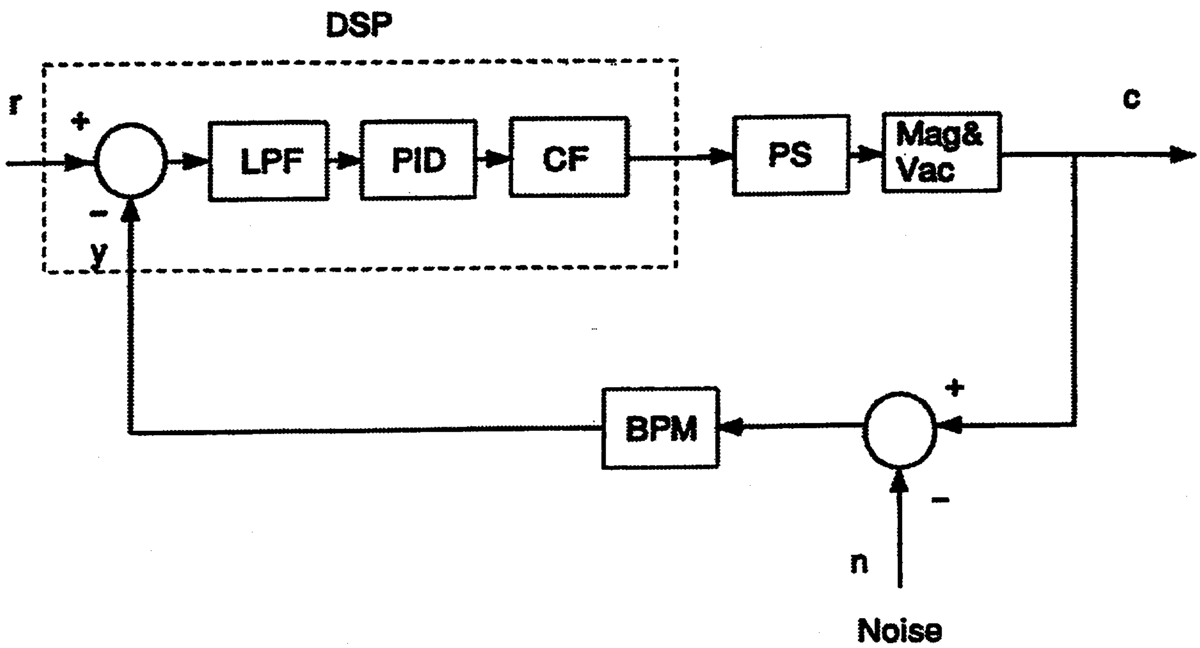

Fig. 1. Simplified block diagram of the APS diagnostic closed loop feedback system

The controller has four main constraints on its design. These are both time-domain and frequency-domain stipulations that should be achieved to allow the closed loop feedback to perform its task. These requirements are:

1. The noise transfer function attenuation should be no less than -12 dB at 20 Hz,
2. The maximum noise transfer function should be less than 2 dB,
3. The system should operate near critical damping to improve time response,
4. The controller gains should be limited to prevent power supply saturation and clipping, which would lead to stability concerns.

The results of the procedure to define a parameter region that satisfies the four main constraints defined by the three gain factors of the PID are shown in Figure 2 [2] where the digital implementation of the PID transfer function is [3]

$$G_c(Z) = K_p + K_i T/Z\text{-}1 + K_d (Z\text{-}1)/TZ$$

and the sampling period is 0.25 ms [2].

The acceptable parameter graph shown in Figure 2 indicates a theoretical practical range of PID parameters to achieve the necessary time and system responses [2]. Various complications can arise in practice that cannot be predicted or compensated. There are numerous areas where these problems can arise; for example the compensation filter designed to cancel the effect of the vacuum chamber, magnet and beam position monitor may perform adequately but not cancel perfectly over the frequency spectrum desired. This in turn would adjust the acceptable ranges of values shown in Figure 2. Other obstacles include temperature, eddy currents, a power supply transfer function that is not truly unity and numerous others which have not been factored into the model. To do so would either be impossible to calculate or would complicate the model to a point of complexity that would make it unusable. Thus simulations may not be entirely accurate, and experiments need to be done to prove the legitimacy of the claims made.

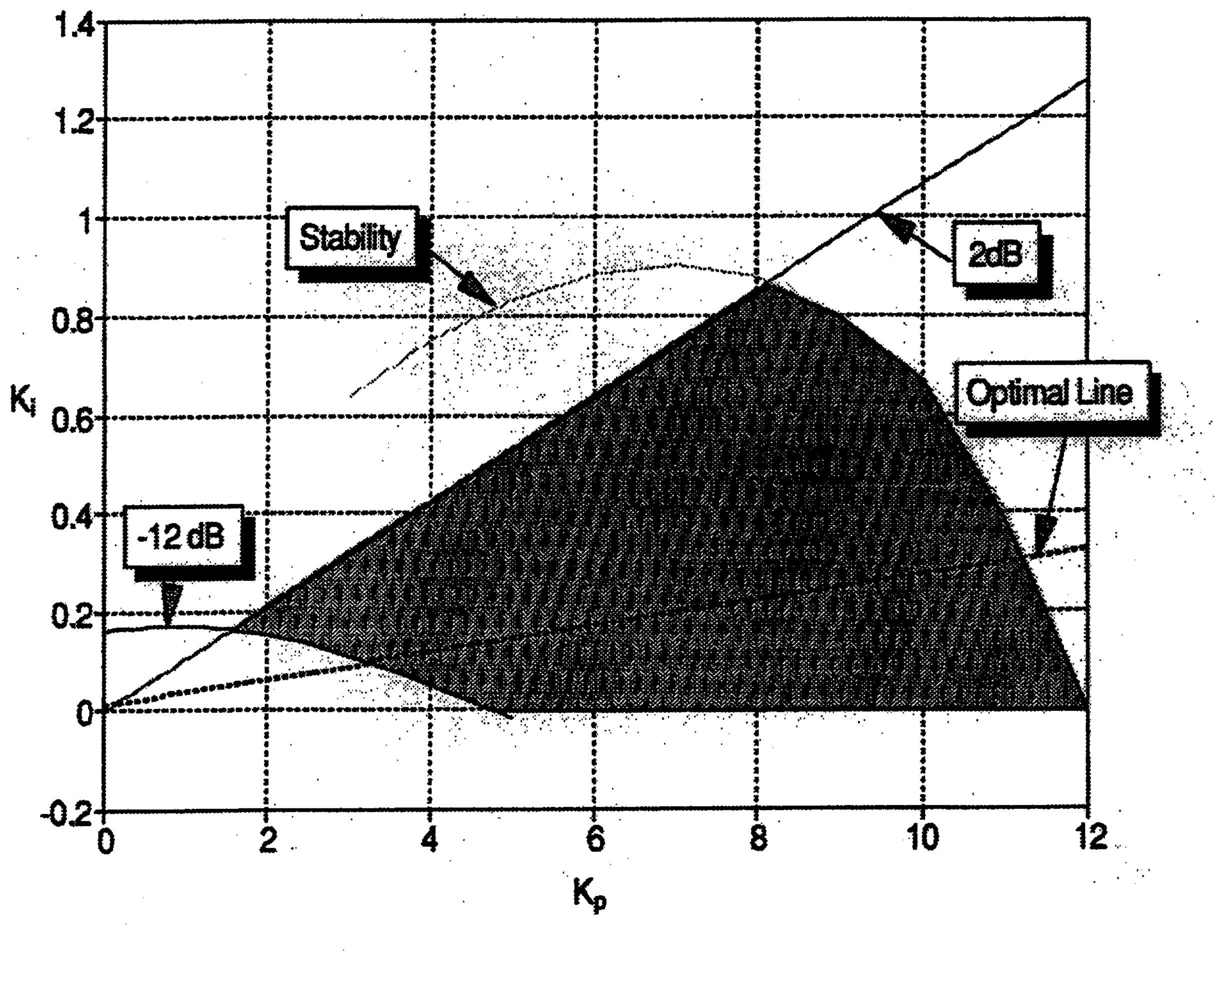

Fig. 2. Acceptable range of parameters for Ki and Kp that satisfy all conditions listed: 20 HZ noise attenuation, maximum noise allowed; system time response and system stability

III EXPERIMENTAL RESULTS

Experiments were conducted using the block diagram shown in Figure 3, the blocks performing the functions of closed loop control. This was implemented using the test setup shown in Figure 4. The DSP completes several tasks in the system. It reduces the bandwidth of the feedback with a low pass filter. This is necessary because the conversion of the analog to digital signals requires the reduced feedback. The low pass filter is anti-aliasing to limit and attenuate the high frequency components. The DSP performs the control portion of the closed loop scheme with a digital PID. It uses the backward rectangular integration digital PID to accomplish this. The final duty of the processor is to compensate for the magnet/vacuum chamber effect in the system.

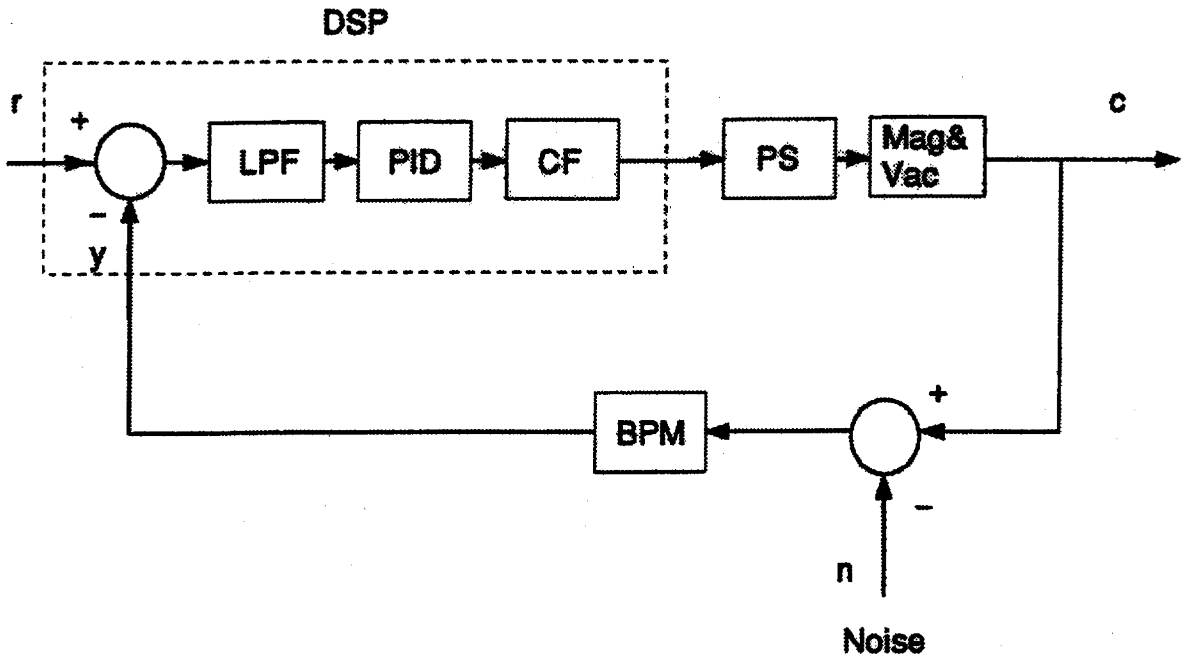

Fig. 3. Block diagram of the actual closed loop system of the global orbit control

Positioned after the DSP is a digital to analog converter on the output. This signal is transmitted to the analog power supply, which for lab experimentation was a Kepco model BOP 20-20M 400 W device. This is a bipolar operational power supply/amplifier capable of being driven by voltage or current and can source or sink 20 A with a voltage capability of 20 V.

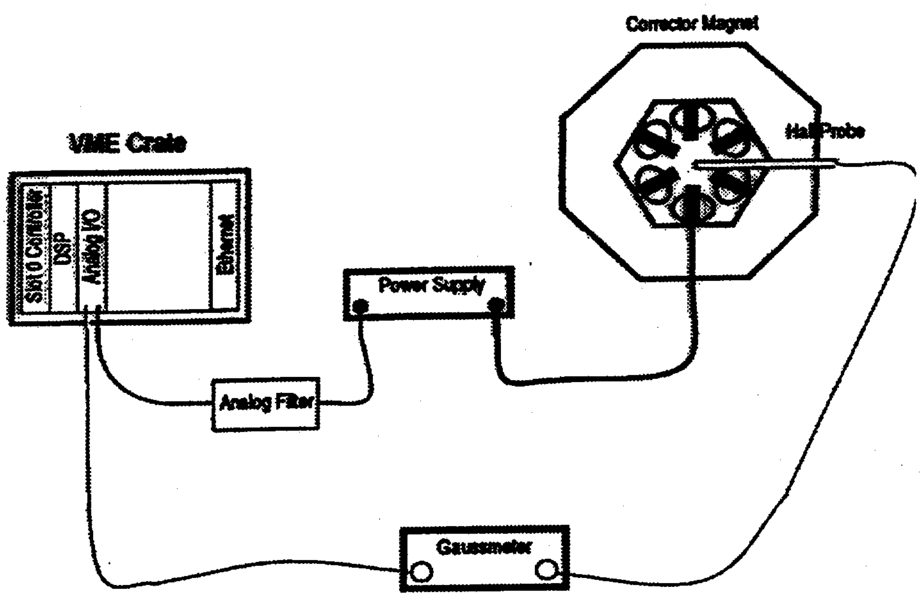

Fig. 4. DSP global orbit feedback test setup

The magnetic field is measured by a Hall effect probe which for purposes of experimentation performs the function of beam position monitor. The beam position is directly affected by the corrected magnet's field, so adequate position simulation is achieved using a probe, together with a gaussmeter, as feedback transducers. This combination measures the magnetic field intensity that reaches the probe through the aluminum alloy vacuum chamber. Eddy currents are formed in this chamber and cause a distortion in the field. This distortion is not static but is dynamically changing, causing difficulty in predicting its transfer function contribution.

The first task performed was the determination of the transfer function for the compensation filter which is in essence the inverse of the magnet/vacuum chamber transfer function. A system analyzer, Hewlett Packard model 3563A, was used to determine this. The analyzer can calculate a digital or analog transfer function for a network by sourcing a test voltage into it and measuring the output. The user may choose the number of poles and zeros in the fit so that a trade-off may be made between fit quality and the complexity and delay required by the filter. The more poles and zeros added the more calculations are needed by the processor, and in general a compromise solution must be accepted.

The first fit contains four poles and three zeros and is depicted in Table 1. The fit is a very good one, and by adding additional poles and/or zeros the response is not significantly improved. The nature of the magnet would lead one to believe that the magnet/vacuum chamber frequency response would have low pass tendencies, and the data and curve fit support this. Since the number of poles is greater than the number of zeros, the net effect is a 20 dB/decade rolloff per unmatched pole for high frequencies.

A different fit contains four poles and two zeros, as is shown in Table 2. This fit does not appear to be so good as the previous one, but the inverted filter does have less gain than in the first case. This is an advantage because the more gain produced by the DSP, the larger the outgoing signal therefore requiring a higher voltage from the power supply, a disadvantage.

The four pole and two zero case becomes a two pole and four zero function when inverted and thus has much gain. Most of this gain appears at high frequencies where there is significant attenuation, but some is placed low enough in the spectrum that it forces the power supply to limit its output. Hence the four pole and three zero curve fit appears to be the better of the two and is the one used for the remainder of this work. Other curve-fit pole/zero combinations were examined, but little was gained by increasing the filter order and much was lost when the order was decreased.

Figure 5 shows the frequency response of both the measured magnet/vacuum chamber and the four pole and two zero digital filter. To cancel the effects of this low pass response of the magnet/vacuum chamber, the digital model filter is inverted to force the product of the two transfer functions toward unity. The ideal case would be unity over the entire frequency spectrum, but since this is extremely difficult frequencies under approximately 1kHz were targeted. That is, more effort went into the design of the compensation filter below 1kHz so that the best compensation would be accomplished in this region. It is for this reason that the number of poles or zeros that are nondominating at high frequency are not critical to the compensation filter design. That design was low order and its performance was adequate with a minimum amount of time delay.

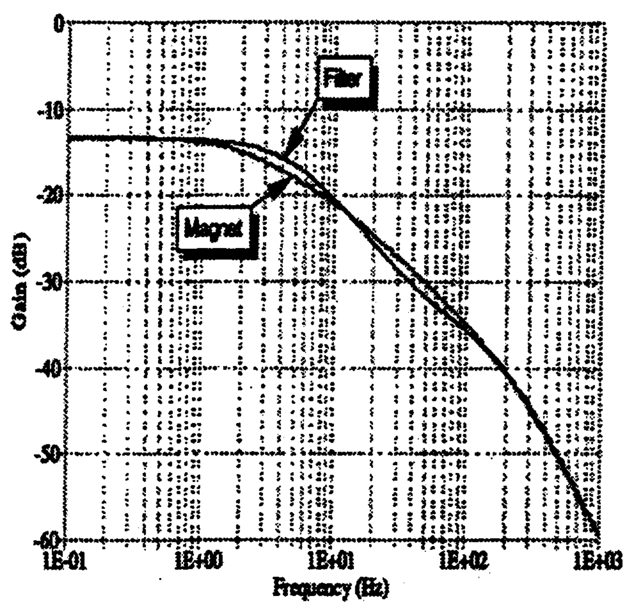

Fig. 5. Frequency response of the digital filter and the actual magent and vacuum chamber

The DSP was programmed with numerous sets of PID values to measure and record the responses. After trying sets of control parameters that ranged from very low to very high gain, it was observed that any set that had even a modest gain caused the power supply to saturate. One method to alleviate such saturation is the addition of a single pole, low-pass analog filter that helps remove high frequency components which are actually amplified in the controller due to the high-pass compensation filter. For it to truly be a compensation filter, its transfer function had to be the inverse of that for the low pass magnet/vacuum chamber. To cancel the magnet/vacuum chamber transfer function, the compensation filter had then to amplify high frequencies, and approached infinite gain as frequency increased. It could not realistically do this, although it may possess an extremely high amplification in this particular range. This is the justification for the analog reconstruction filter (ARF), which prevents the power supply from limiting due to its inability to switch at that rate.

IV PID INFLUENCE ON ATTENUATION OF NOISE TRANSFER FUNCTION

The first constraint considered is the noise transfer function attenuation, which is a means to measure the ability of the control system to reduce the effects of feedback noise. It is the goal of the control system to have the largest attenuation possible at 20 Hz, below which most vibrational noise is predicted to be present. Figure 6 indicates the PID values used for the experiment.

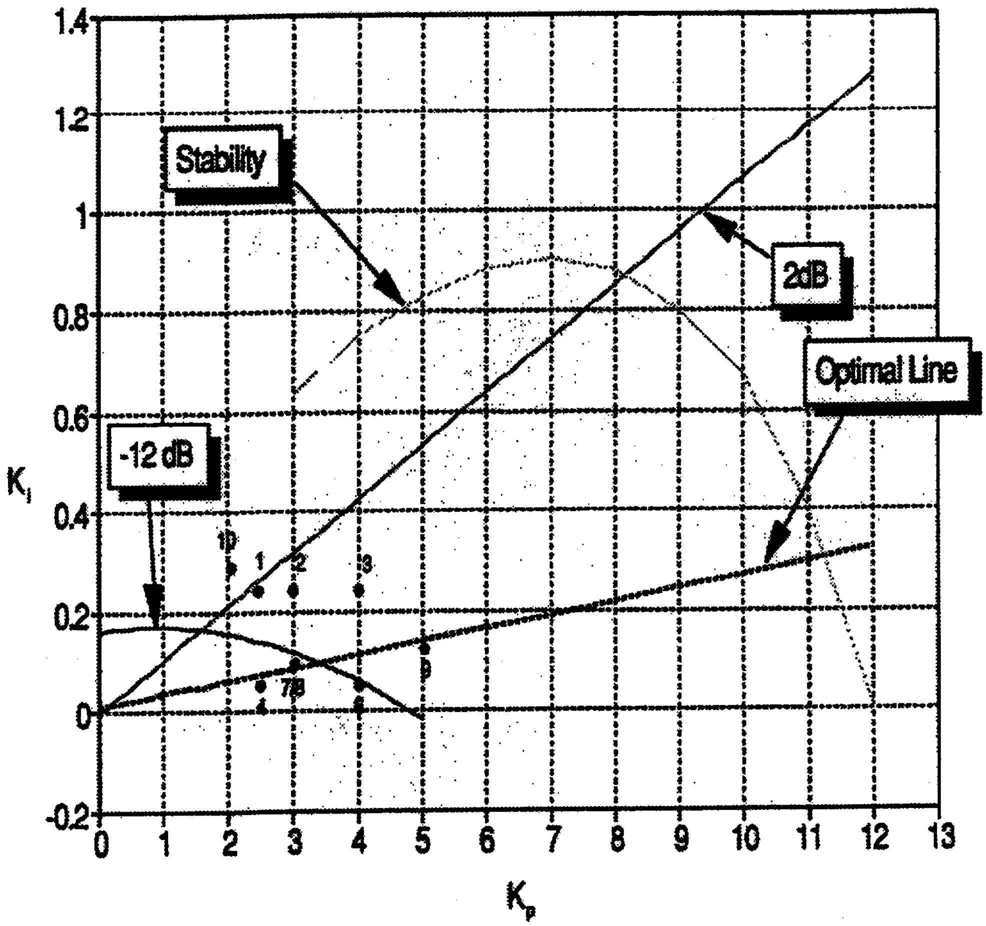

Fig. 6. Points selected to measure system response in the frequency and time domain

The first group of points to be examined are the K_i=0.25 with K_d=0 and K_p ranging from 2.5 to 4.0. This allows analysis of K_p and its impact on noise feedback. Each group of PID sets is subjected to three ARF cutoff frequencies. Figure 7 displays the frequency response of the noise transfer function with K_p=2.5, K_i=0.25, and K_d=0. The three traces shown are taken using this PID with different filter cutoffs. The high frequency noise in the system is readily apparent using a 1000 Hz ARF. The lower frequency range under 1 Hz seems erratic apparently due to the DSP's compensation filter amplifying high frequency noise, and the ARF's cutoff not being low enough to adjust it. The 20 Hz attenuation does not appear to be shaped by the filter. Table 3 shows the frequency results of a series of experiments for ranges of K_p, K_i, and K_d. From this table, all of the first grouping is seen as approximately -18 dB. The second group deviates from a barely unacceptable -11 dB to an adequate -14 dB, depending upon K_p.

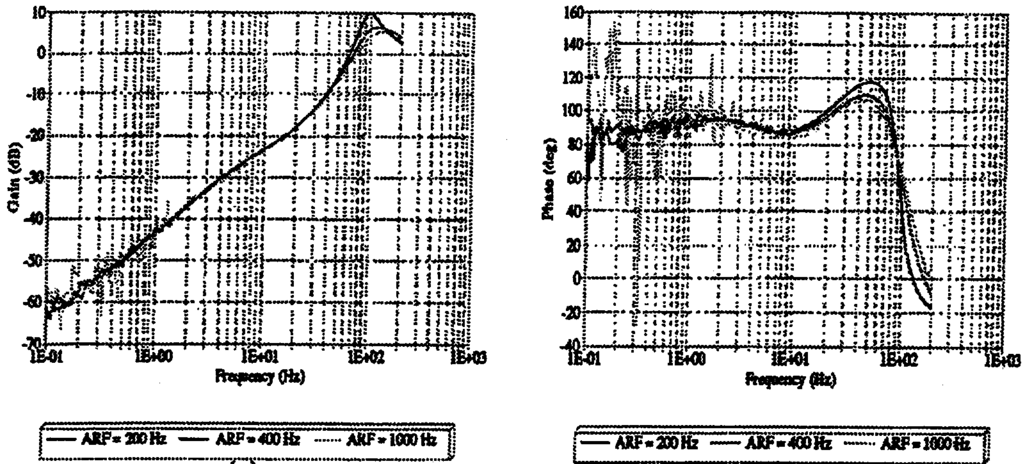

(a) (b)

Frequency response of the noise transfer function with K_p=2.5, K_i=0.25, and K_d=0.

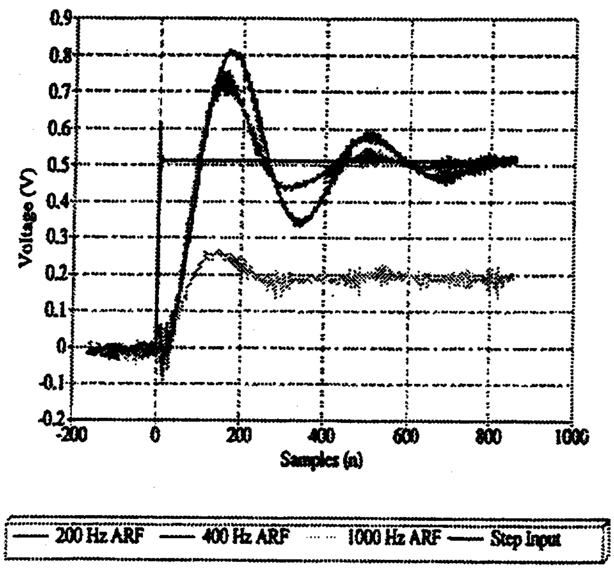

Step response of K_p = 2.5, K_i = 0.25, and K_d = 0.

Figure 7

Since the proportional and integrative constants are most important, they have received the principal attention. It should be noted that the best attenuation occurred with the last point (K_p=2, K_i=0.3, K_d=0) and the worst with one of the first (K_p=2.5, K_i=0.05, K_p=0). It appears from the table that the integration constant has the most influence on noise control. It is the integration portion of the PID that gives the closed loop control a pole at DC, so this should be reflected in the value chosen for K_i. It should also be noted that a higher gain assigned to K_p seems to diminish noise. As the gain of G is increased, the effective noise transfer function is decreased. Since the gain of G is directly proportional to the controller proportional constant, any increase in controller gain will reduce noise output.

V PID INFLUENCE ON SYSTEM TIME RESPONSE

In a subsequent experiment a step function was applied to the system and by saving the data points stored in the signal analyzer the output response was measured and recorded. The step response was taken using the points of Figure 6 and, as is shown in Figure 7b, the effect of the ARF is apparent. The added pole increases the overshoot and influences the rise time. Because it was necessary to reduce the input signal for the 1000 Hz ARF, it is difficult to compare that case with the other two. It does appear that the 1000 Hz condition has a smaller overshoot than the others, with the 200, 400 and 1000 Hz ARF cases having 60%, 40% and 25% overshoot respectively.

Examination of Figures 8 and 9 reveals more about the system time response. The value of K_i has been reduced from 0.25 to 0.05 and this has changed the step response significantly. When the proportional constant K_p is 2.5 the system response is slightly underdamped with minimal overshoot. However as K_p is increased the overshoot and settling time escalate. This indicates that decreasing the integrator gain K_i while increasing the proportional gain does not diminish the overshoot.

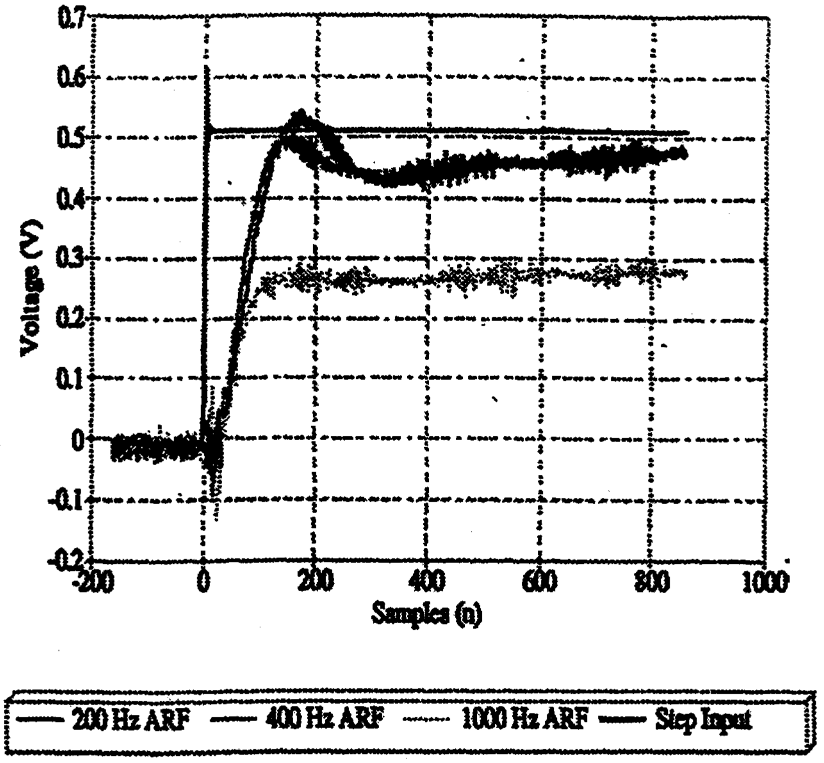

Figure 8 Step response of $K_p = 2.5$, $K_i = 0.05$, and $K_d = 0$.

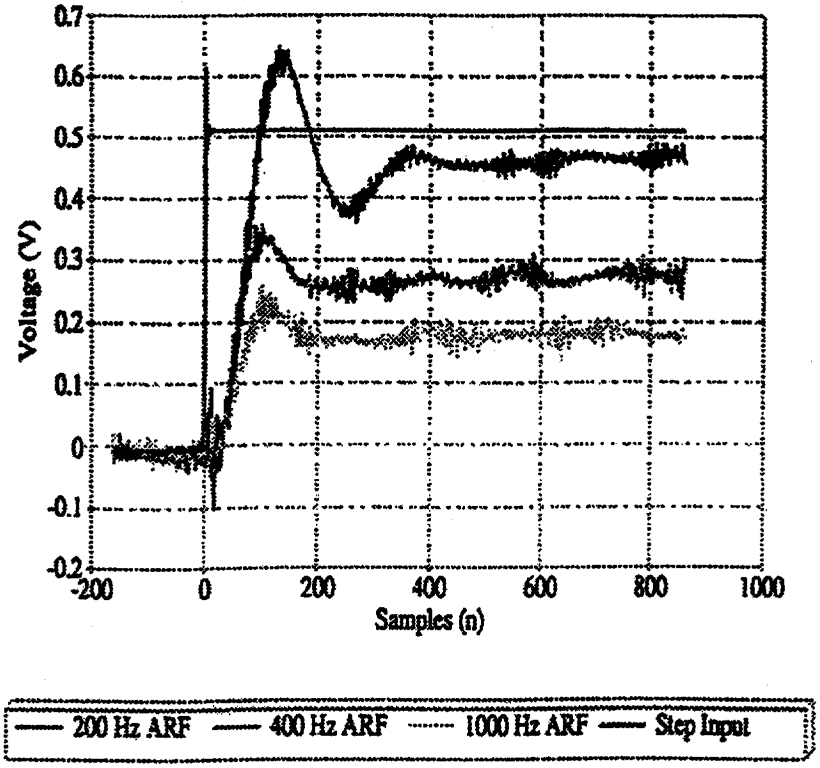

Figure 9 Step response of $K_p = 4$, $K_i = 0.05$, and $K_d = 0$.

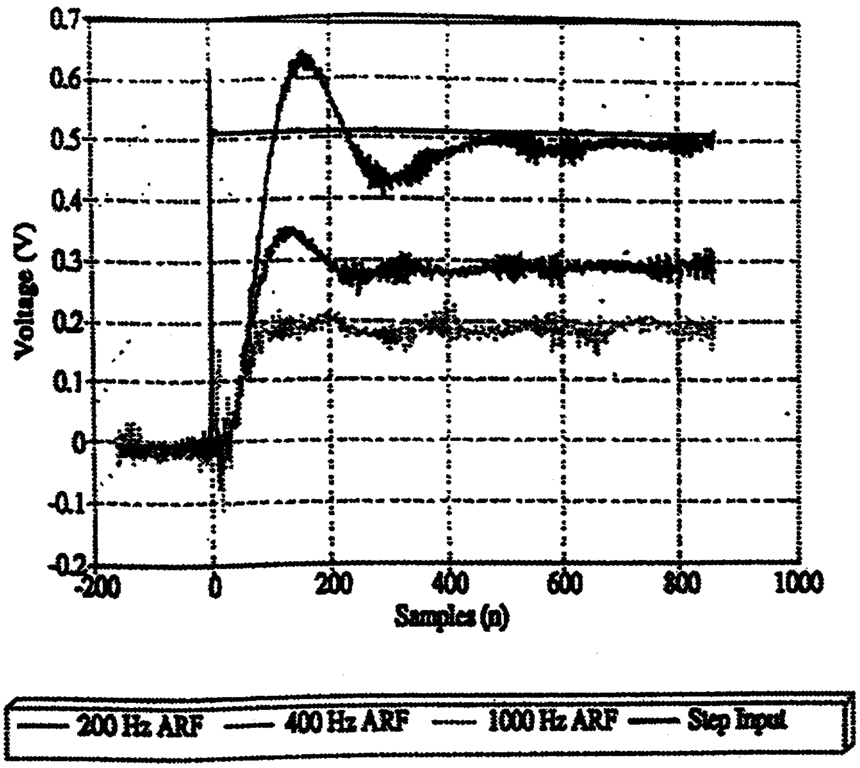

Figure 10 Step response of $K_p = 3$, $K_i = 0.1$, and $K_d = 0$.

What happens if derivative control is utilized to help offset integration control is indicated in Figures 10 and 11. Using values of $K_p = 3$, $K_i = 0.1$ and K_d either 0, no control, or 1, one can see that without derivative control the 200 Hz ARF overshoot is approximately 25%, but with it the overshoot decreases to 20%, not a huge benefit but an advance nonetheless. With the 400 Hz ARF a more pronounced change in that response appears, being converted from underdamped to critically damped.

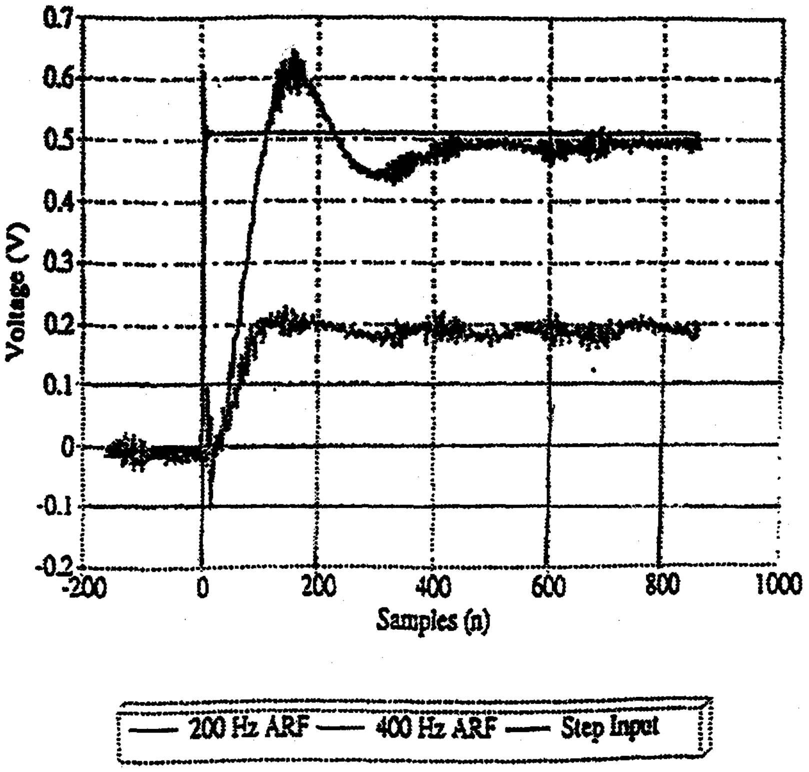

Figure 11 Step response of $K_p = 3$, $K_i = 0.1$, and $K_d = 1$.

Table 1 Filter coefficients of the first compensation filter design.

$a_0 = 9.41294 \times 10^{-4}$	$b_1 = -1.87602$
$a_1 = 1.30784 \times 10^{-2}$	$b_2 = 7.05058 \times 10^{-1}$
$a_2 = -8.52438 \times 10^{-4}$	$b_3 = 3.46876 \times 10^{-1}$
$a_3 = -1.11792 \times 10^{-2}$	$b_4 = -1.74941 \times 10^{-1}$

$$H(z) = \frac{\Sigma a_k z^{-k}}{1 + \Sigma b_k z^{-k}}$$

Table 2 Filter coefficients of the second compensation filter design.

$a_0 = 1.36339 \times 10^{-2}$	$b_1 = -2.87463$
$a_1 = -7.6617 \times 10^{-4}$	$b_2 = 7.16485 \times 10^{-1}$
$a_2 = -1.091 \times 10^{-2}$	$b_3 = 3.25489 \times 10^{-1}$
	$b_4 = -1.66391 \times 10^{-1}$

$$H(z) = \frac{\Sigma a_k z^{-k}}{1 + \Sigma b_k z^{-k}}$$

Table 3 Noise transfer function measurements including maximum gain, frequency of maximum gain, 20 Hz attenuation, and phase margin.

K_p	K_i	K_d	ARF	sweep	G_{pk}	f_{pk}	G@20Hz	P.M.
1	0.25	0	200	200	10	98	-18	114
			400	200	6	119	-18	99
			1000	200	5	119	-19	89
3	0.25	0	200	200	9	108	-18	109
			400	200	6	130	-18	93
			1000	200				
4	0.25	0	200	200	9	143	-19	104
			400	200				
			1000	200				
1	0.05	0	200	500	4	125	-11	60
			400	500	4	172	-11	51
			1000	500	3	202	-11	45
3	0.05	0	200	500				
			400	500				
			1 000	500				
4	0.05	0	200	500	7	147	-14	75
			400	500	6	202	-14	64
			1000	500	6	202	-14	59
3	0.1	0	200	200	6	124	-13	75
			400	200	6	124	-13	63
			1000	200	7	113	-14	61
3	0.1	1	200	200	6	124	-13	70
			400	200	6	124	-14	59
			1000	200	6	119	-13	64
5	0.125	0	200	200	7	173	-16	88
		0	200	200	7	173	-16	87
		1	200	200	7	165	-16	88
2	0.4	0	200	200	8	81	-22	152
		0	200	200	8	81	-22	149
		1	200	200	8	81	-22	148

The most important questions are: Did the simulation tests support the data taken for the actual time responses? Does the optimal line live up to its name? If the 1 kHz ARF cases are examined and compared the optimal line seems to hold true. However it is important to realize that the 1 kHz ARF case is closest to simulation conditions where there actually was no analog reconstruction filter. The 1 kHz cutoff puts the pole far enough away from the rest of the response that its influence is not great.

The points 4, 7,8 and 9 of Figure 6, nearest the optimal line, display little or no overshoot with the exception of point 9. This point uses a 200 Hz ARF, which has adversely influenced the response to reflect the added pole. Since the other two points on this line approach critical damping, there is no other reason for point 9 to do so. These other points also have excellent rise times and settling times, while if one deviates from this optimal line the effects can be readily seen.

VI CONCLUSION

The APS third generation synchrotron light source requires a global orbit control system that reduces noise effects and stabilizes the beam. This was accomplished through a high speed feedback loop established completely around the accelerator ring. This encompassing control system requires a specific kind of control that balances its needs without compromising any of these demands too heavily.

The need exists for a controller to achieve noise reduction, but other consequences must be considered as well. Such notable system attributes as time response, stability, power supply bandwidth and cost are all concerns of closed loop control and it was necessary to find a control algorithm that met or exceeded the design criteria. The PID controller proved to be the best in this situation, and design centered on the choice of proportional, derivative and integral gains and sampling time

Experimentation has justified the type of controller design, where optimal control was achieved using three-dimensional plots of the performance in question. These plots were constructed with two of the design parameters as variables; they can also be organized into tabular format to facilitate comparison viewing. The performance measures were rated in terms of noise filtering ability, system response time or system stability. This paper has described methods to assess system performance of both time and frequency domain functions.

Performance of noise rejection was measured by plotting the noise transfer function value at a critical frequency for different values of K_p, K_i, and K_d where the values that met or exceeded the criterion for handling injected noise were plotted. The same procedure holds true for determination of maximum gain of the noise transfer function.

The work performed in this paper examined the controller's design effects on the performance of the system with one "BPM" and one corrector. A better experiment could be performed using all 40 of the BPMs and correctors. This would be the final proof of the efficiency and capability of the controller design. This experiment would need to be performed on an existing accelerator where a PID loop could be programmed into a global orbit feedback system.

REFERENCES

1. D.Bulfone et al., "Design Considerations for the ELETTRA Beam Position Feedback Systems," European Particle Accelerator Conference, 1990, PP 895-897.

2. J. A. Kirchman and J.P. Bobis, "Preliminary Design of the APS PID Global Orbit Control System," IEEE International Conference on Industrial Electronics, Control and Instrumentation, Maui, Hawaii, Nov. 1993.

3. Benjamin C. Kuo, Automatic Control Systems, Prentice-Hall, Englewood Cliffs, NJ 1991.

FERMILAB-96/069

New SWIC Scanner/Controller System

W. Kissel, B Lublinsky, A. Franck

Fermi National Accelerator Laboratory*

Abstract

Since the early 1970s Segmented Wire Ion Chambers (SWICs) have been used in the Fermilab Switchyard beam transport lines to measure beam positions, emittances and intensities by generating horizontal and vertical beam profiles. Each plane utilizes 48 wires positioned in an electric field to collect beam ionization current. By integrating the current in each wire, two-dimensional beam profiles can be generated. The Switchyard SWICs as originally engineered are geographically multiplexed. There are 26 physical devices and only 7 scanners/controllers; therefore all profile data cannot be collected in one beam cycle. In the Main Ring era short cycle times (~10 seconds) lessened the impact of this limitation. Tevatron era operation of Switchyard uses both SWICs and BPMs. SWICs will show profile information and have a much lower beam intensity threshold as compared to BPMs. This is an advantage during tune-up at low intensity especially when dealing with cryogenic transport lines. Long Tevatron cycle times (~60 seconds) coupled with the need to sequence various SWICs and their gains on a given controller, have been seen to increase the required startup/tuning time, necessitating the upgrade presented here.

Hardware Implementation.

The new SWIC controller/scanner system (figure 1) has been designed to ease the limitations noted in the current one. Each SWIC has an associated rack-mounted controller chassis which contains an integrator channel per wire, timing and control circuitry, a dedicated 80186 CPU and an ARCNET communication coprocessor. The controller is capable of taking and storing successive beam snapshots under various timing and gain settings. Triggers for data collection are derived from the Tevatron real-time clock or from direct processor commands. The SWIC position, in or out of beam, and the associated high voltage supply are interlocked and controlled by the chassis as well. All of the SWIC controllers reside on an ARCNET LAN and deliver data to a front end processor (68060 residing in the VME chassis). This front end processor does data organization and statistical processing before the final data is presented to the control consoles for operator analysis. New cabling combined with a new scanner/processor system will allow simultaneous data collection from all SWICs. Existing "Automated Tuning" programs will be able to use either BPM or SWIC position information.

Software Implementation.

Software support for the SWIC system will include programming SWIC scanners to do certain measurements, store the results of the measurements and calculations performed on them in various buffers, and provide convenient user access to these buffers. A major problem in this system is that of somehow tying together multiple various pieces of data belonging to simultaneous measurements.

The most convenient way to do this, in our opinion, is to use the Finite State Machine (FSM) paradigm. If we apply this paradigm then we can view the VME system as running 26 FSMs (one per actual SWIC processor). These FSMs are comprised of states, each corresponding to one measurement. Every state is set up independently, defining the type of measurement (fast spill, slow spill, early in spill, late in spill, beam-off background, etc.) it represents, integration time, integration time prescale, gain, calculations to be performed on the collected data, where this data should be written (associated buffers), etc. Each FSM is also settable, allowing the user to define which states are to be executed and in which order.

* Operated by the Universities Research Association under contract with the U.S. Department of Energy

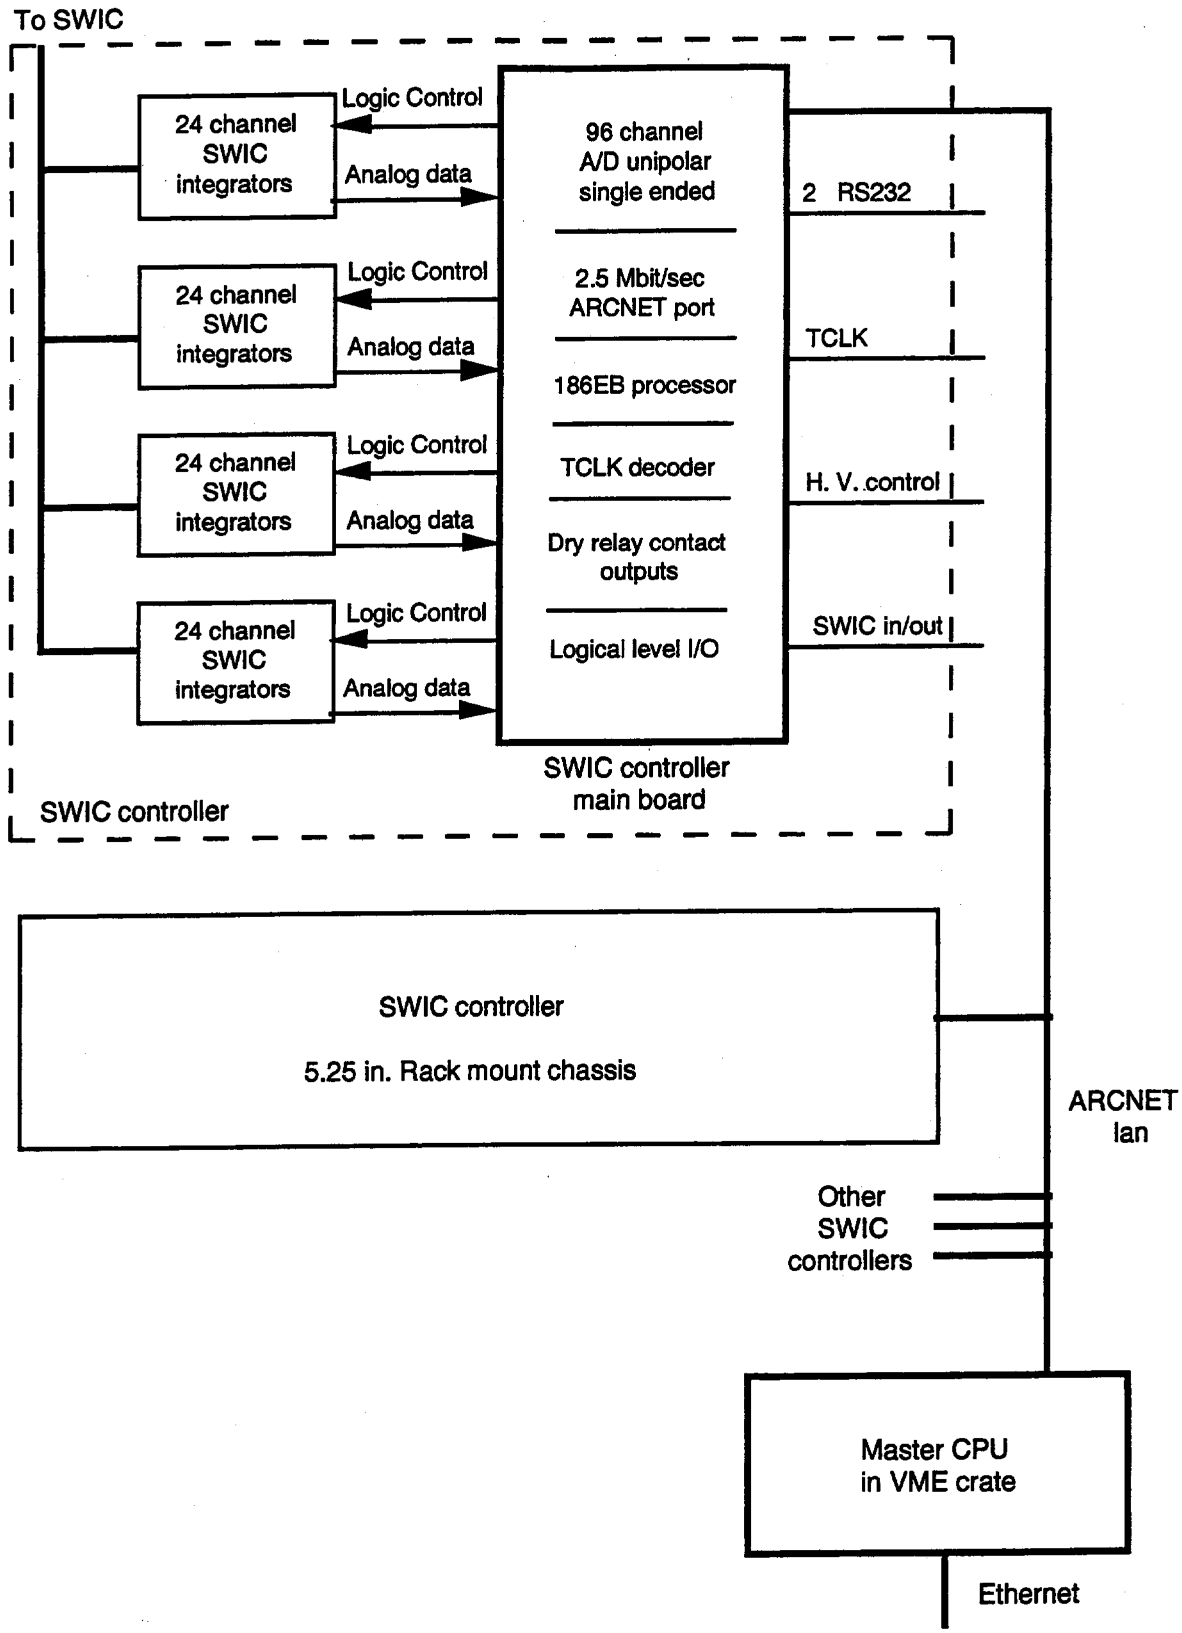

Fig 1. SWIC control and data retrival system.

SWIC software is architecturally similar to that of Tevatron cryogenics [1]. Both are multiprocessor systems, consisting of a master communication CPU and multiple data collection slaves. The communications between master and slaves in both systems are based on the ARCNET link. This allows us to use without any changes that ARCNET protocol created for the cryogenic system. This protocol provides prioritized task to task communications and has proved to be reliable and easy to use. Reflective memory support created for the cryogenic control system, the virtual I/O bus (VIOB), can also be used for support of the bulk of the SWIC devices that are read and set periodically. The only extension for SWICs is data from the integrators, which is returned from the slaves to the master synchronously and bypasses the VIOB.

Summary and Conclusions.

This project has proved to be an interesting application of the reuse of software. Although the SWIC beam diagnostics discussed here differ markedly from the industrial process control of cryogenics, nonetheless distributed system operation, at a level below Fermilab ACNET, has been easily ported between them. The new software required is in the form of device drivers for the integrator boards and mathematical treatment of the profiles.

References

1. B. Lublinsky, J. Firebaugh, J. Smolucha, New Tevatron Cryogenic Control System, "The Proceedings of the 1993 IEEE Particle Accelerator Conference", V3, p1817.
2. B. Lublinsky. Shell software for Smart subsystems with front-end capabilities. Nuclear Instruments and Methods in Physics Research A 352 (1994) pp403 - 406

Use of an INGRES database to implement the beam parameter management at GANIL.

Gillette Pascal, Lécorché Eric, Lermine Patrick, Maugeais Clotilde (Computers Group)
Leboucher Christian, Moscatello Marie-Hélène, Pain Pascal (Parameters Group)

GANIL BP5027 14021 CAEN France

e-mail : lecorche@ganexp.in2p3.fr

Since the beginning of the operation of the new Ganil control system in February 1993, the relational database management system (RDBMS) Ingres has been more and more widely used. The most significant application relying on the RDBMS is the new beam parameter management which has been entirely redesigned. It has been operational since the end of the machine shutdown in July this year.

After a short recall of the use of Ingres inside the control system, the paper first explains how the parameter management is organized. Then it shows the database implementation and how the physical aspects of the Ganil tuning have been integrated in such an environment.

Lastly, the paper tries to draw some outlines by showing more generally what advantages the use of the relational database brought in this context.

1. THE DATABASE INSIDE THE CONTROL SYSTEM.

1.1 The control system renewal.

From 1989 to December 1992 the Ganil control system was completely upgraded by replacing the previous and obsolete control system by a new architecture. The new system [1] relies on an Ethernet network onto which several kinds of processors are connected. VAX/VMS computers are used either for development purposes or as a control server at the real-time level. Each operator console consists of the logical association of an X-terminal and a workstation onto which is plugged a dial-box to implement the knob function. More than 2500 pieces of equipment have to be handled through front-end processors which are in VME or CAMAC crates and running the VAXELN operating system.

Most of the software is written in the ADA language and graphic developments rely on the Motif standard.

1.2 The relational database management system inside the control system.

Although it was a new technology in our environment, it was decided rather early in the project to build the data management for the new control system on a Relational Database Management System (RDBMS). Local considerations led us to choose the Ingres RDBMS for this and when the new system went into operation in February 1993, the RDBMS was already involved in many fields (ref. [2]) :

- The first important use was the equipment data management : software and hardware addresses, device scaling information, units etc. Files are extracted from the database and downloaded into the front-end crates or installed in shared memories in the workstations. For this application, access from the real-time level to the database is strictly limited to the update of trace flags for debugging and statistics.

- An other major part of the system built upon the database was the alarm logging system. Alarms are displayed on an X-terminal and VT consoles and at the same time they can be (if flagged) stored in the database for a short or long term period. This functionality is achieved by two processes communicating through mailboxes. A first process is in charge of the real-time level and display and it only interfaces the database after changes concerned with the alarm configuration; the second one directly stores the alarms or their acknowledgments in the database. This application constituted the first real-time application related to the RDBMS.

- Due to the capabilities brought by the RDBMS, it was also decided to manage basic operator menus in another Ingres database where tables are organized in a recursive way. Menus are loaded into the process memory when starting its execution and the coupling with the database is rather weak.

- An off-line application is in charge of the daily operation journaling. The database is filled by the operators without any interaction with the real-time level.

- The first version of the so-called "beam parameter database" BDPARAM was designed to store and manage the results from the off-line program PARAM which calculates the theoretical values of most of the machine devices according to the ion beam to be accelerated. The database was used as a gateway between the PARAM program and the on-line tuning programs which use files extracted from the database.

Graphic applications specific to the database are performed under the Ingres/Windows4GL environment. For the control software, integration of the ADA/SQL access is done through a preprocessor invoked before compiling.

1.3 First evolution.

The RDBMS was first applied to many aspects of the control system but apart from the control level. Due to the rising emergence of the relational technology in accelerator controls at that time and our lack of experience with such a system, we wanted to follow a very careful approach in this domain and as previously described direct links between the Ingres database and the on-line level were very restricted.

After the experience gained during the first year of operation with the new control system, the database was progressively integrated into the control level as we got a deeper knowledge of the RDBMS and as people came to appreciate the benefits brought by the RDBMS. The capabilities and features induced by such an approach are actually much more important than the interfacing response times found in most cases.

The most important development following this path is the new design for the beam parameter management which is now entirely built upon the BDPARAM database. The database has been widely extended and is accessed by on-line control applications as explained in this contribution.

2. THE BEAM PARAMETER DATABASE.

2.1 Ganil operation.

The Ganil can accelerate many different ion beams characterized by the beam specifications consisting of the particle to be accelerated, the ion charge at the source output and after the stripper, the beam energy, the RF frequency etc. As it is quite a complicated machine with three cyclotrons in cascade, the machine tuning involves sending the beam into several machine configurations, each of them with several possible beam optics for tuning or measurement. After the beam production and acceleration, it is adjusted and analyzed through a spectrometer before being sent to the experimental switchyard.

2.2 Machine description.

The design of the beam parameter database was first required to provide a complete description of the machine. So, all the entities associated to the beam parameters i.e. the pieces of equipment, beam characteristics etc. have been ordered according to their position along the beam path. Furthermore, these entities have been collected into objects sorted within an object-oriented approach by defining classes of entities: basic pieces of equipment (quadrupoles, dipoles, steerers, strippers, NMR probes etc.); more complex devices (RF systems, Magnetic spectrometers etc.); global machine components (cyclotrons etc.); beam properties (beam characteristics). Objects of each class have a fixed number of entities referred to by a predefined type. (For example, quadrupole objects have two entities of the "current" and "gradient" types.)

2.3 Database implementation.

The beam parameter database BDPARAM has to store all the values involved in the machine control for any part of the machine as managed by the object entity description. The aim is to provide on-line programs relying on the database and able to manage any machine configuration then to set all the pieces of equipment to the corresponding settings through the front-end crates; these on-line programs constitute the "PARAMETERS" application family. Beam parameters come either from the theoretical calculation (PARAM program) or previous settings. Settings for the beam lines are dynamically calculated by the PARAMETERS applications, as described below. The database is schematically represented in figure 1 and actually consists of more than 100 tables.

Each ion beam is specified by its BEAM_ID in the BEAM_DESCRIPTION table which is attached to a theoretical beam parameter set issued from the existing off-line program PARAM. This theoretical parameter set stored into the THEORETICAL_VALUES table contains the beam characteristics (the beam specifications and other beam parameters such as the magnetic rigidity etc.) and all the theoretical parameters independent of the beam course and optics (i.e. the cyclotrons).

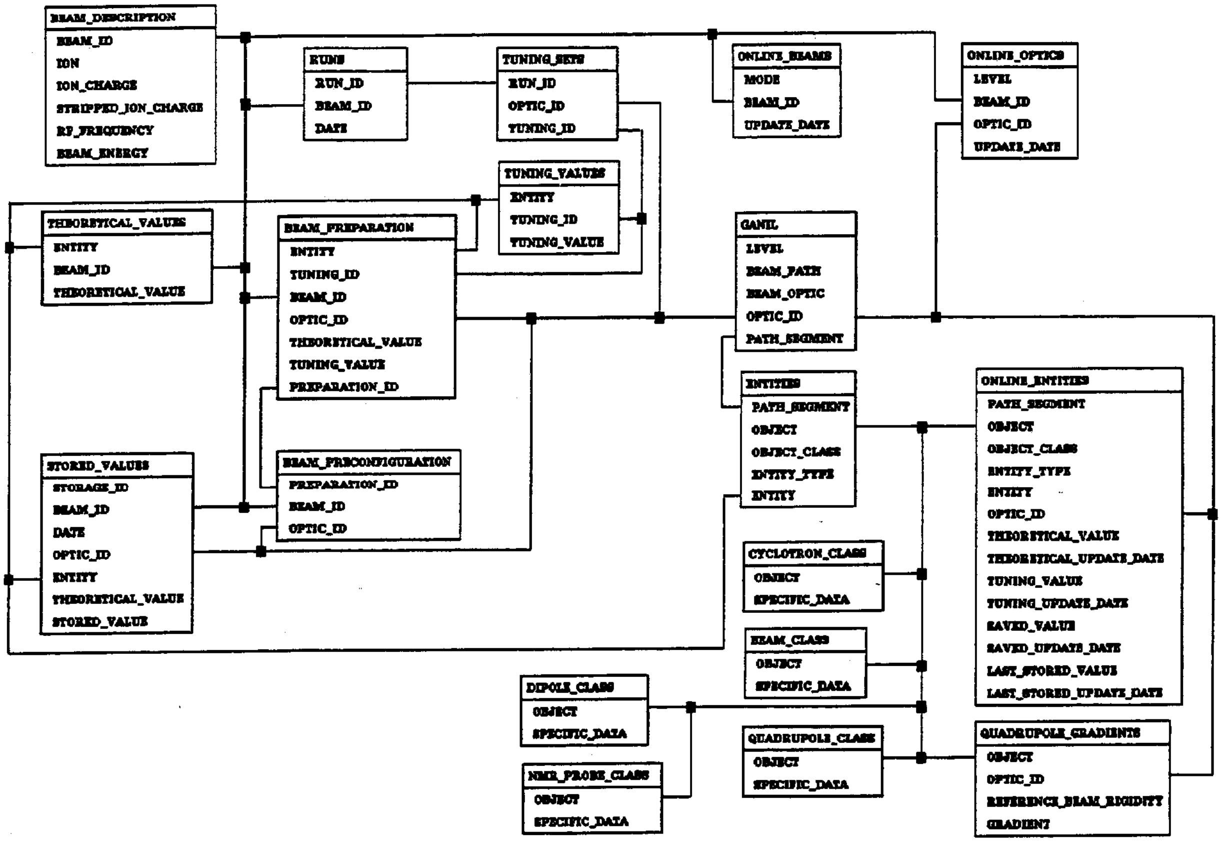

Fig 1 : Schematic overview of the beam parameter database.

The various tuning configurations achievable for both the GANIL machine and the experimental areas are listed into the GANIL table according to the BEAM_PATH with the corresponding BEAM_OPTIC optics configuration for each of them. Each beam path / beam optics couple belongs to an operational level of the Ganil facility; these levels are the "Machine", "Alpha spectrometer", "Medium Energy Output", "Experimental area distribution" and "Experimental Rooms". Any BEAM_PATH is divided into intrinsic segments named PATH_SEGMENT, each segment is defined as an indivisible element for the beam tuning and can belong to several beam paths.

The ENTITIES table then provides the complete description of the entities according to the object decomposition as described previously. Their adherence to the beam paths is seen through the PATH_SEGMENT attribute of the GANIL table. Object classes needing specific data for their complete definition involve the creation of particular tables (BEAM_CLASS, DIPOLE_CLASS, QUADRUPOLE_CLASS etc.). Also the QUADRUPOLE_GRADIENTS table had to be added defining the gradient value for each beam optic in which the quadrupole is involved according to the beam rigidity chosen as a reference.

Short term parameters set can be stored into the STORED_VALUES table ; they can be archived for a long term storage into the TUNING_VALUES table seen through the TUNING_SETS and RUNS tables.

Some tables are updated dynamically during the machine operation. First of all, the ONLINE_BEAMS table describes the beams currently produced both for the on-line and off-line modes, as one of the injector cyclotrons can deliver a "local" beam independently from the beam given to the physicists. The ONLINE_OPTICS table therefore specifies the beam path and beam optic currently applied to any of the machine levels. Lastly, the ONLINE_ENTITIES table is an on-line extension of the ENTITIES table used to keep dynamic object values.

3. MANAGEMENT OF THE BEAM PARAMETERS.

3.1 General overview of the data Flow diagram.

The beam parameter management is performed at two levels. Firstly, the off-line level consists of the parameter calculation by the PARAM program and the preparation of the machine operation. Then the on-line Motif programs PARAMETERS directly interface with the database and set pieces of equipment to the appropriate values obtained from the database management.

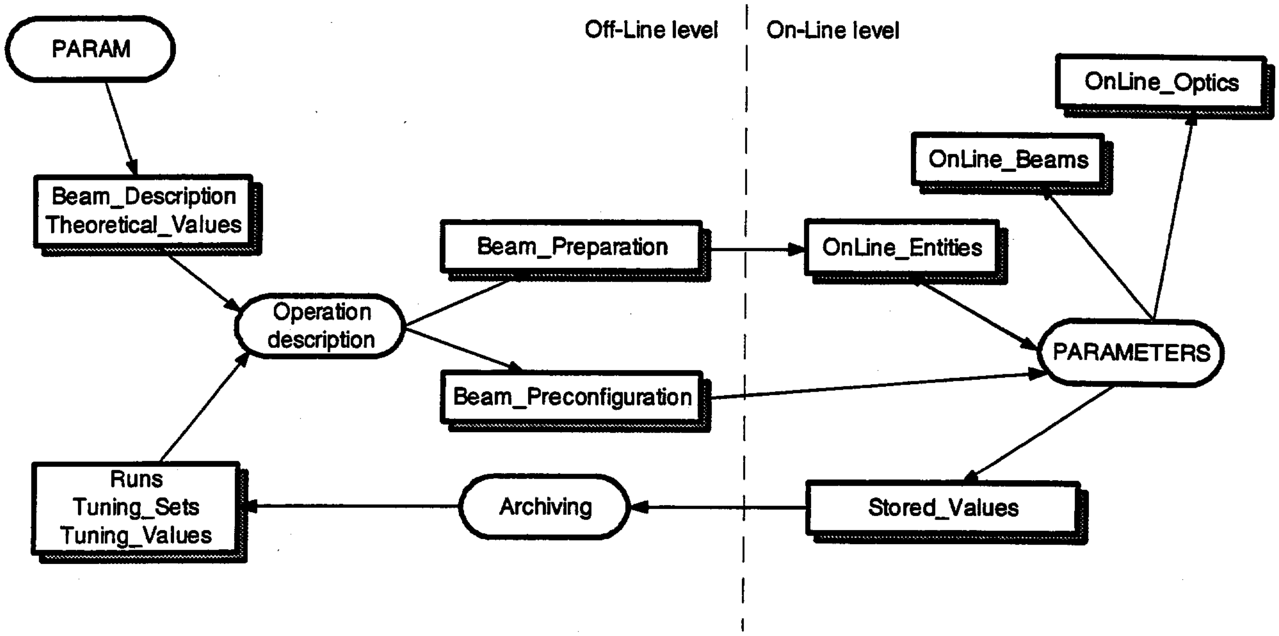

Fig 2 : Data flow diagram overview.

3.2 Off-line and on-line calculation.

The former FORTRAN off-line program PARAM calculates the theoretical values for any given beam specified by the ion to be accelerated, the energy required and the beam characteristics. SQL procedures load into the THEORETICAL_VALUES table only the theoretical values independent of the beam path and optic and not yet integrated inside the on-line program. All the beams known in the database are referenced in the BEAM_DESCRIPTION table are also updated from the PARAM program. The parameters concerning the beam lines, that is to say the optics parameters, are calculated on-line, all the data necessary for the calculation being in the database. The data flow diagram overview is presented in Figure 2.

3.3 Beam Preparation.

As the Ganil machine can accelerate a large variety of beams, it is necessary to prepare the beams which are going to be produced. So SQL operation description tools allow people to define beam parameter sets in the BEAM_PREPARATION table: first a "*basic set*" consisting only of the cyclotrons and beam parameters is generated from the theoretical tables independent of the beam path and optics to be applied later; then other "*dedicated sets*" attached to beam path / beam optic couples can be added, mixing the theoretical values issued from PARAM and archived values from the RUNS, TUNING_SETS and TUNING_VALUES tables.

The BEAM_PRECONFIGURATION table gives the on-line programs the basic information for the parameters set to be associated with a dedicated tuning mode.

3.4 On-line management.

The on-line management is done by the PARAMETERS applications allowing the operator to perform all the tuning phases concerning the machine configuration. First, when a beam has been chosen to be operated from the BEAM_PRECONFIGURATION table, the ONLINE_ENTITIES table is updated from the BEAM_PREPARATION table either for the theoretical or tuning values found in the basic set attached to the beam. The ONLINE_ENTITIES table is also updated after any optic change at any level of the machine : values can again be moved from the preparation tables (with the dedicated sets) but for the beam lines they also can be calculated dynamically. From this table, pieces of equipment can then be set to the theoretical, stored or archived values. The PARAMETERS applications also operate automatic beam rigidity adjustment either from current or stored values. Lastly, the scope of any action can be restricted graphically by the operator.

The current status of the whole machine is kept in the ONLINE_BEAMS and ONLINE_OPTICS tables.

The operator can also save any part of the machine in the ONLINE_ENTITIES table for a later recall or comparison purposes. Short-term storage is done into the STORED_VALUES table consisting of all the entities belonging to the beam path. They can be restored globally into the ONLINE_ENTITIES table or archived for further use using SQL archiving tools.

3.5 Software implementation.

The PARAMETERS applications are written in ADA and profit from the client-server architecture. Interaction with the operator is achieved through the Motif interface. The XRT/table graphic widget has been used for the data presentation included in the on-line applications.

Due to the internal multitasking capability of the language, ADA packages had to be first developed to interface to the Ingres database. They implement an internal database server for the application process to prevent concurrent multitasking access to the database, as the DBMS does not know the multitasking ADA specific features.

Most of the Ingres features have been used in the database design including database procedures (integrity rules, complex data management etc.), database rules (data integrity, coherence etc.) and database events raised on specific operation phases (beam or optics changes, modifications of the experimental area configuration).

In fact the database is duplicated on two different machines. The first version is used off-line level and comprises the complete database including the empty on-line tables to be used as backup if needed. The second one is an extraction from the first, consisting only of the data directly involved in the machine operation. Specific replication tools allow people to work on the off-line database and then to transfer their data in a secure way to the on-line level.

4. PRESENT STATUS AND FUTURE ENHANCEMENTS.

4.1 Status.

The whole beam parameter management system has been set into operation since the machine startup in July this year. From the operation point of view, it has already brought a higher reliability and more flexibility in the machine tunings; some specific experiments have benefited from the new capabilities offered. Also the group in charge of the beam parameters can now better cope with the increasing number of beams to manage, in a faster and more flexible manner than before. Moreover, the analysis of tuning parameters and statistics can be done more easily.

Some work is in progress to improve the whole system and new possibilities and features are being smoothly introduced, such as the propagation of the most important events dealing with the operation through asynchronous events in the database. Except when starting the on-line applications, response times are not a problem in this system. This aspect has been examined by looking carefully at the physical structure of the tables and the way ADA/SQL programming is done. We also use the statistics collected by the RDBMS to improve the optimizer efficiency.

Finally, access to the BDPARAM database through the ADA/SQL packages is beginning to be used in applications which had previously been unable to access any beam parameter values.

4.2 Extensions.

What has been done just constitutes the first and basic phase of the beam parameter management. In a second step, new tools or applications will be developed taking benefit of the database. At the off-line level, some tools for data analysis have to be defined as well as graphic tools for the beam preparation in a more convenient way. The on-line programs will progressively integrate the present off-line calculation: new object classes will therefore need to be created to follow this trend, the final objective being to be able to provide an autonomous on-line beam parameter calculation program.

4.3 Benefits of the database.

The beam parameter database is a typical illustration of the impact of a RDBMS in an accelerator control system. Looking at the benefits offered by such a system and its intrinsic properties (reliability, coherence, maintainability, methodology, global approach etc.), we can assume that the design of any control system should incorporate a database management system to provide all the required efficiency.

Considering that the challenge was not so obvious six years ago, the final conclusion could be that the use of Ingres inside the Ganil control system has proved to be a real success.

REFERENCES.

[1] L. David, E. Lécorché, C. Maugeais, M. Ulrich : "Controlling the GANIL from an Operator Workstation", Proc. Int. Conf. EPAC'94, London, Great-Britain 1994

[2] E. Lécorché, P. Lermine : "Database management in the new GANIL control system", Int. Conf. ICALEPCS'93, Berlin, Germany 1993

The Advanced Photon Source Injection Timing System*

Frank R. Lenkszus and Robert Laird
Argonne National Laboratory

ABSTRACT

The Advanced Photon Source consists of five accelerators. The injection timing system provides the signals required to cause a bunch emitted from the electron gun to navigate through intermediate accelerators to a specific bucket (1 out of 1296) within the storage ring. Two linacs and a positron accumulator ring operate at 60Hz while a booster synchrotron ramps and injects into the storage ring at 2Hz. The distributed, modular VME/VXI-based injection timing system is controlled by two EPICS-based input/output controllers (IOCs). Over 40 VME/VXI cards have been developed to implement the system. Card types range from 352MHz VXI timing modules to VME-based fiber optic fanouts and logic translators/drivers. All timing is distributed with fiber optics. Timing references are derived directly from the machine low-level rf at 9.77MHz and 351.9MHz. The timing references provide triggers to programmable delay generators. Three grades of timing are provided. Precision timing is derived from commercial digital delay generators, intermediate precision timing is obtained from VXI 8-channel digital delay generators which provide timing with 25ns peak-to-peak jitter, and modest precision timing is provided by the APS event system. The timing system is fully integrated into the APS EPICS-based control system.

INTRODUCTION

The Advanced Photon Source consists of five accelerators: an electron linac, a positron linac, a positron accumulator ring, a booster synchrotron, and the positron storage ring itself. The injection timing system is required to generate all the timing pulses (with the exception of the linacs which have their own timing subsystem) needed to cause bunches to be emitted from the electron gun, accumulated in the positron accumulator ring (PAR), accelerated by the booster synchrotron to 7GeV, and finally injected into selected buckets of the storage ring.

In order to position the bunches in the rf buckets of the three circular machines, the injection/extraction timing pulses for a machine must be synchronized to that machine's rf frequency. The booster and storage ring use the same rf frequency of 351.9MHz. The PAR's fundamental rf frequency is 1/36 that of the storage ring or 9.77MHz. The PAR also uses the 12th harmonic of this frequency to compress its stored bunch. There are 36 rf buckets in the PAR. Since the circumference of the PAR is exactly one wavelength of its fundamental, the 36 rf buckets are accomplished by shifting the phase of the fundamental in 10-degree steps. The booster, which is one-third the circumference of the storage ring, has 432 rf buckets while the storage ring has 1296 rf buckets.

The timing system uses the PAR fundamental rf frequency and the booster/storage ring rf frequency as reference inputs. These inputs are processed by digital circuits to generate timing references which, in turn, are used to trigger digital delay generators. The digital delay generators produce the actual timing pulses. Two grades of digital delay generators are used. For those timing pulses requiring very precise timing, Stanford Research Systems DG535s [1] are used. The DG535s are used for applications such as kicker timing and linac modulator timing. Less demanding requirements are fulfilled with in-house-designed VXI 8-channel digital delay generators. The VXI delay generators digitally count a clock to generate programmed delays. A peak-to-peak jitter of 25ns is obtained when the VXI generator's internal 40MHz clock is used. Nearly all timing signals are transmitted over fiber optic links to their destination.

The Advanced Photon source event system is also an integral part of the injection timing system, but since it is described in detail in another paper [2], it will not be discussed here.

INJECTION TIMING CYCLE

Figure 1 illustrates the APS injection cycle. During storage ring injection, the cycle repeats at a 2Hz rate. Nominal operation requires 24 linac pulses to be accumulated in the PAR at a 60Hz rate. This is followed by a 100ms period in which the PAR bunch is damped and compressed. At about 50 ms after the last bunch is injected into the PAR, the booster synchrotron is commanded to ramp. When the booster reaches the injection energy (~450MeV) about 30 ms after ramping is started, the bunch is ejected out of the PAR and into the booster. The booster then accelerates the injected bunch to 7GeV during the next 220ms, at which point the bunch is extracted out of the booster and into the storage ring.

The injection timing system is required to put the bunch into any of the storage ring's 1296 buckets. This is accomplished by shifting timing reference pulse generation to correspond with the desired bucket. The module that generates timing references for PAR injection uses a PAR rf reference from the rf bucket phase shifter as its internal logic clock. Thus, the generated timing references track the phase-shifted rf. The timing references for PAR extraction, booster injection and extraction, and storage ring injection are shifted digitally via counters that are preset to the desired bucket number and clocked by the booster/storage ring low-level rf.

*Work supported by U.S. Department of Energy, Office of Basic Energy Sciences, under Contract No. W-31-109-ENG-38

PRIMARY TIMING TRIGGERS

Figure 2 is a greatly simplified representation of the timing system. It consists of two input/output controllers (IOCs) separated by about 300m. One IOC, iocinjtime, is located in the injection control room, while the other, iocmtime, is located in the main control room. Iocinjtime controls timing for linac triggers, PAR injection/extraction, booster injection, and the associated beam transfer lines. Iocmtime controls booster extraction, storage ring injection, and the associated beam transfer line. Each IOC is VME based with an attached VXI crate and GPIB bus. The IOC's processors run the EPICS control system software under vxWorks. All timing parameters are controlled through EPICS graphical control screens.

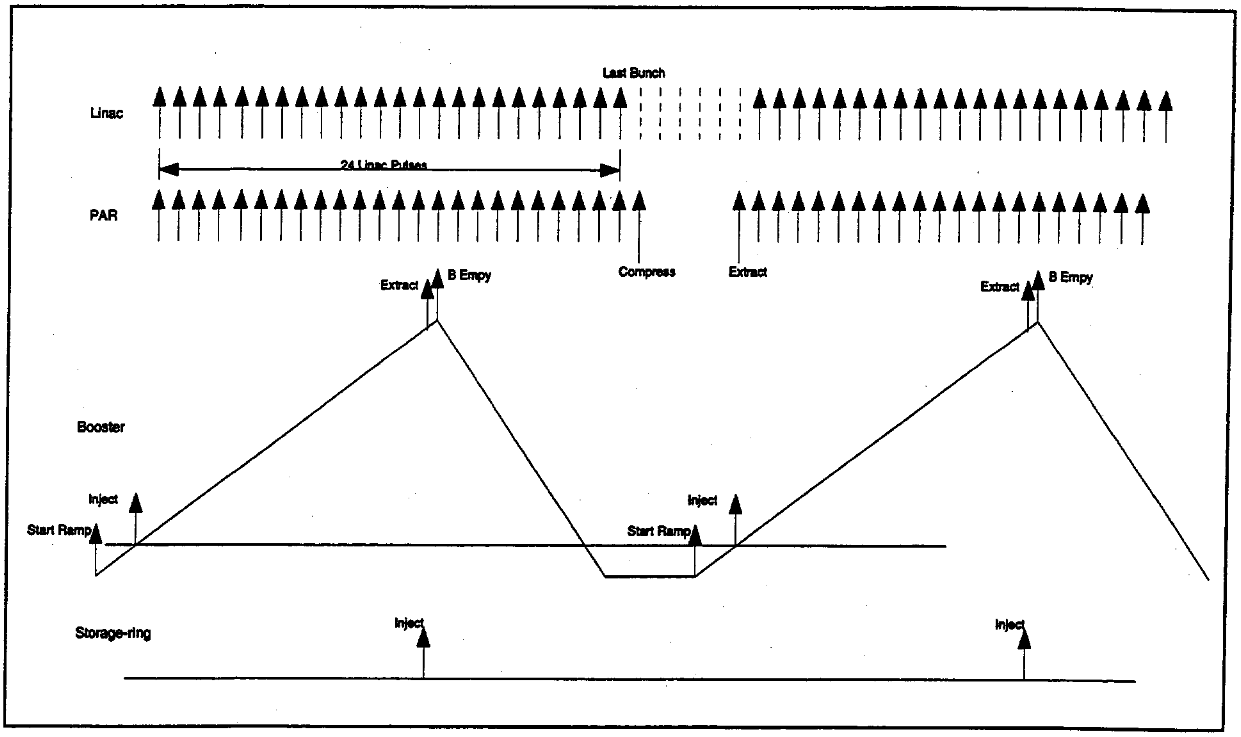

Figure 1 The APS Injection Timing Cycle

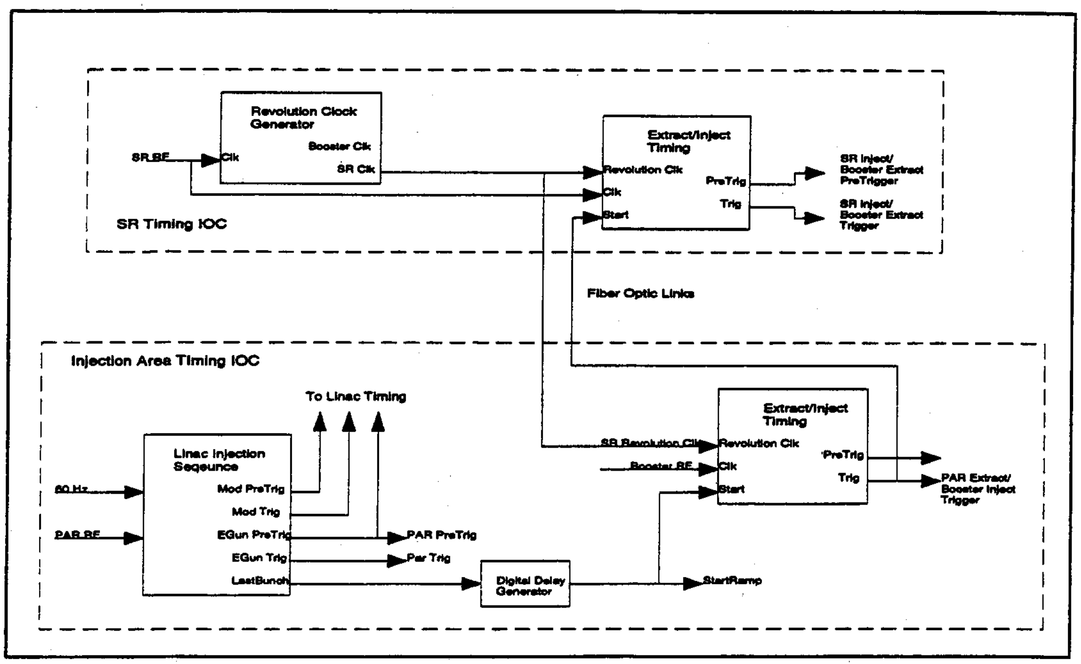

Figure 2 Injection Timing System Block Diagram

Each IOC generates many timing signals to control kickers, septa, diagnostics, etc. Most of these signals are generated by digital delay generators that are triggered by timing fiducials. The timing fiducials are derived from three primary references: the 60Hz line frequency, the PAR fundamental frequency, and the booster/storage ring rf frequency. The booster/storage ring rf frequency is synchronously counted by the Revolution Clock Generator module in iocmtime to produce revolution clocks for the booster synchrotron and the storage ring.

Timing fiducials for linac to PAR are generated by a Linac Injection Sequence module which is described in greater detail below. The fiducials are: linac modulator pretrigger (to time modulator charging), linac modulator trigger, linac electron gun (egun) pretrigger, PAR pretrigger (controls pulsed supply PFN charging), PAR trigger (times pulsed magnet firing) and last-bunch which provides a timing reference for things that happen during the PAR compress and extract period.

An injection cycle begins with the generation of egun pretrigger/trigger timing references. A pretrigger occurs 1/60th of a second before a trigger. The pretrigger triggers a set of digital delay generators that commands the PAR pulsed magnet power supplies to charge their PFNs. The egun pretrigger also sets up the linac timing subsystem to trigger the egun on the next modulator trigger. The egun trigger timing fiducial triggers a set of digital delay generators that generates PAR pulsed magnet power supply discharge commands. The egun pretrigger and trigger timing fiducials are synchronized to the first PAR rf cycle zero crossing following a 60Hz line zero crossing. Thus, the egun pretrigger and trigger are synchronized to both the PAR rf and the 60Hz line frequency.

A last bunch timing fiducial is generated coincident with the last egun trigger pulse in the sequence. This signal provides a timing reference for things that are to occur after the PAR is filled, such as turning on the PAR 12th harmonic rf to compress the bunch and commanding the PAR extraction and booster injection pulsed magnet power supplies to charge their PFNs. The last bunch is also used to trigger a digital delay generator which generates a start ramp timing fiducial. Start ramp initiates the booster synchrotron ramp cycle.

When a start ramp occurs, the timing system counts a preset number of storage ring revolution clocks to determine when to generate the PAR extract/booster inject timing trigger. This timing reference triggers a number of digital delay generators which produce discharge commands to the PAR extraction and booster injection pulsed magnet power supplies. This timing reference also triggers the PAR-to-booster beam transfer line timing generators.

The booster inject timing trigger is transmitted via a fiber optic link to iocmtime which controls booster extraction. The booster injection trigger causes iocmtime's Extract/Inject Timing module to begin counting storage ring revolutions. This module is loaded with the number of revolutions required for the booster to accelerate the bunch to 7 GeV. The acceleration period is approximately 220 ms. About 40 ms before booster extraction, the Extract/Inject Timing module generates the booster extract/storage ring inject pretrigger. This timing reference triggers digital delay generators that command the booster extraction and storage ring injection pulsed magnets to charge their PFNs. At 350 microseconds before booster extraction, the module generates the trigger timing reference which triggers digital delay generators that produce discharge commands for the booster extraction and storage ring injection pulsed magnets. Beam transfer line diagnostics are also timed relative to this trigger signal. Upon bunch extraction, the booster is commanded to ramp down in preparation for the next cycle.

Different storage ring buckets are filled by shifting the timing references generated by the Extract/Inject timing modules in the two timing IOCs. Each of these modules counts storage ring rf cycles relative to the revolution clock to select different buckets.

LINAC INJECTION SEQUENCE

Figure 3 show a greatly simplified block diagram of the Linac Injection Sequence trigger (LIST) module. This VXI module generates timing references that are used to control linac triggering and PAR injection. It accepts the 60Hz line frequency and PAR rf frequency as inputs. The internal logic is completely synchronous and uses the PAR rf input as the clock.

The module uses two 30-bit synchronized shift registers to generate two sequences of pretrigger and trigger timing references. One sequence times the linac modulators while the other times egun triggers and PAR injection. Any 30-bit pattern may be specified for either the modulator or egun trigger sequences. The bit patterns are shifted out at the 60Hz rate and repeated at a 2Hz rate. Each shift register has an associated sequence length counter which determines the end of a sequence. The pretrigger and trigger outputs are 1/60th of a second apart and always occur in pairs.

All outputs from the LIST module are synchronized to the first PAR rf cycle zero crossing following a 60Hz line zero crossing. Although not shown in the simplified diagram, the LIST module contains a delay counter which can shift the timing of all outputs relative to the 60Hz line by an even number of PAR rf cycle to a maximum of 128k.

EXTRACT/INJECT TIMING MODULE

Figure 4 shows a simplified block diagram of the VXI module which generates the pretrigger and trigger timing references for PAR extraction, booster injection/extraction, and storage ring injection. One of these modules is located in both iocinjtime and iocmtime.

The module's internal logic is completely synchronous to minimize timing errors due to logic propagation delay drifts. The module uses the 352MHz booster/storage ring rf as its internal clock. A 16-bit presettable counter counts rf cycles after the occurrence of the storage ring revolution clock. This counter is preloaded with the desired bunch number. When the bunch number is reached, the counter generates a bunch select signal which is connected to the count inputs of two additional presettable counters—

the turns counter and the pretrigger delay counter. These two counters begin counting bunch selects upon occurrence of the start input. The turns counter is preloaded with the number of turns to occur between start and the pretrigger timing reference out. The pretrigger delay counter is preloaded the number of turns desired between the generation of pretrigger and trigger. Load bunch is an external input which is driven by a signal that is timed (via a digital delay generator) to update the bunch select register with a new value (written by the IOC) between timing cycles.

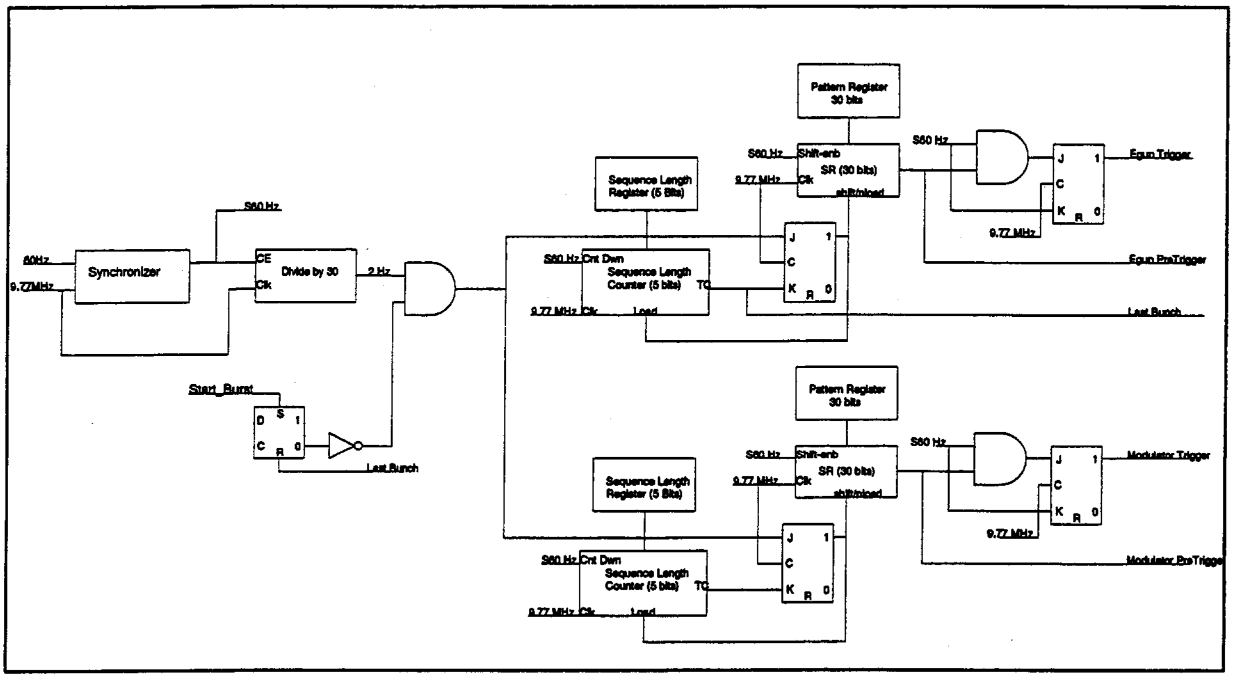

Figure 3 The Linac Injection Sequence Trigger Module Block Diagram

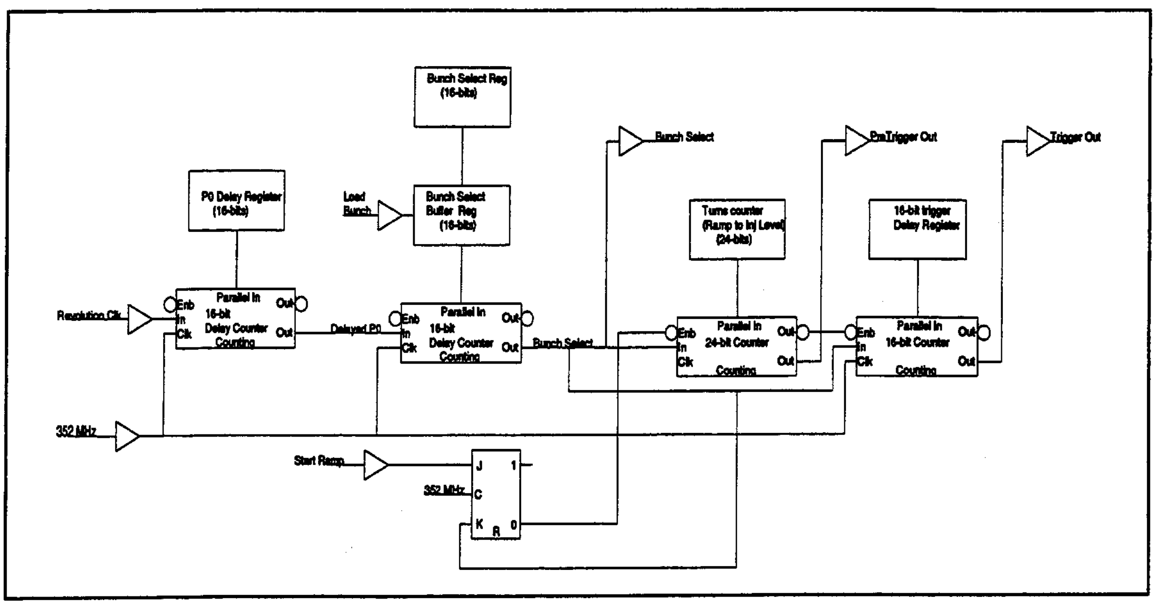

Figure 4 The Extract/Inject Timing Module Block Diagram

The operator need only set things such as desired storage ring bucket number and booster ramp time. The IOC processors calculate the proper numbers to load into the Extract/Inject Timing modules. As mentioned previously, the module uses the 352MHz low-level rf as a clock for its synchronous logic. We have, however, successfully tested the module with input clocks of up to 490MHz.

REVOLUTION CLOCK

The revolution clock VXI module produces revolution clocks for the booster and the storage ring. This module uses synchronous counters to divide the 352MHz booster/storage ring rf reference by 432 to obtain the booster clock. The booster clock is in turn divided by three with another synchronous counter to produce the storage ring revolution clock. The module also has an input for the PAR 9.77MHz fundamental rf. This input is obtained from a source ahead of the PAR bucket phase shifter and is used to determine which of the 432 booster and 1296 storage ring buckets will be bucket "0." Upon command from the IOC, the revolution clock counters are synchronously reset at the next 9.77MHz zero crossing. Under normal circumstances, this resyncing of the revolution clocks need only be done at system initialization time.

VXI DIGITAL DELAY GENERATOR

Figure 5 shows a block diagram of the 8-channel VXI digital delay generator developed for general-purpose timing. The eight channels have independent programmable selection of clock source and trigger source. Each channel may be triggered from any of the VXI TTL trigger lines or from one of four front panel trigger inputs. The clock source may be the internal 40MHz clock, the internal clock divided by 2 or 4, or the external clock input.

The counters for each channel are contained in an individual programmable logic device (PLD). Thirty-two bits are available for pulse width/delay determination. The number of bits assigned to width and delay is determined by how the PLD is programmed. We have three standard configurations: 8-bit width x 24-bit delay, 16-bit width x 16-bit delay and 24-bit width x 8-bit delay. A channel's delay and width range depend on which type of PLD is socketed for that channel. The 8-bit width x 24-bit delay is used most frequently. With this configuration, pulse widths of up to 6.4 microseconds and delays of up to 419 milliseconds with a resolution of 25 nanoseconds are possible.

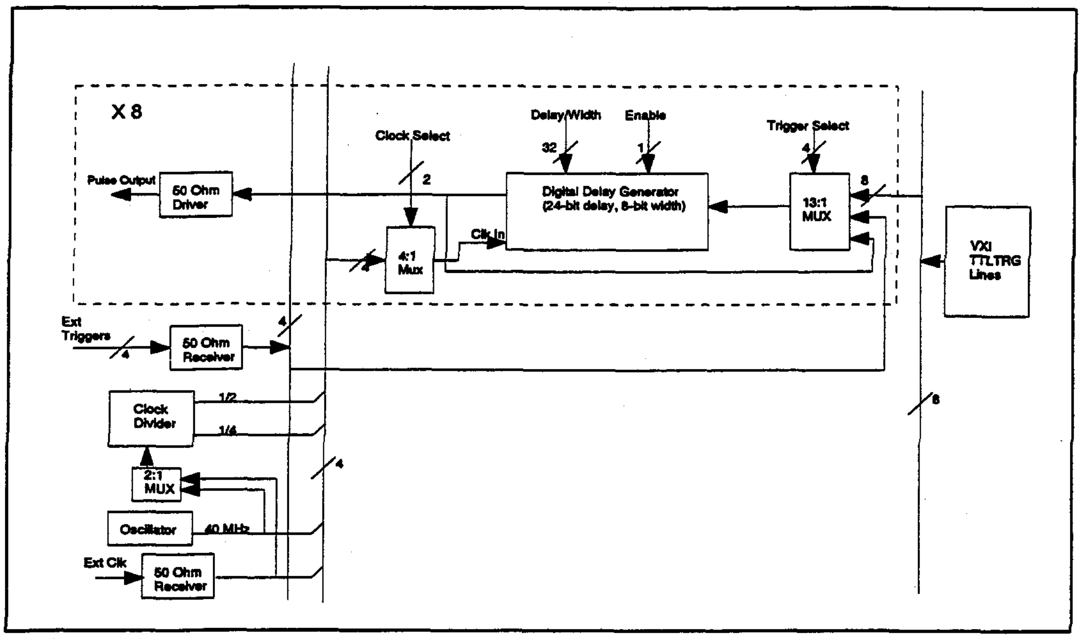

Figure 5 Eight-channel VXI Digital Delay Generator

OTHER MODULES

All together about 40 different VME and VXI modules have been designed to implement the APS injection timing system. Other modules include: VME video sync generator, VME revolution clock counter/trigger, fiber-optic transmitters and receivers, fanouts, and logic-level converters. All the modules are plug and play and are supported by EPICS. VME fanout and logic-level convertor modules use the VME backplane for power only; i.e., they have no programmable features.

CONCLUSION

The timing system has been in place for nearly two years. Although a good deal of effort was expended initially to determine timing requirements, many additional ones arose as the systems began to come on-line. The timing system's modularity made it possible to respond to all these new requirements in a timely fashion.

REFERENCES

[1] Stanford Research Systems, Inc., Sunnyvale, Ca.
[2] F. R. Lenkszus, R. Laird, "The Advanced Photon Source Event System," these proceedings.

Finite State Machine Data Acquisition at Fermilab

D.A. Herrup, B. Lublinsky, J. Wang

Fermi National Accelerator Laboratory*

ABSTRACT

Modern Accelerator control requires the collection and manipulation of large quantities of data in response to accelerator-related operations and timing events. We have implemented a system at Fermilab based on a set of finite state machines which allows state-by-state definition of the data collection parameters and the storage of up to 100 KB of data in an easily accessible set of circular buffers. Initially the system will digitize beam position detectors and provide tune information. We will describe some of the other uses for this system.

INTRODUCTION

One of the important problems in accelerator control concerns the ability to synchronize and collect data describing various processes, and eventually to provide control information for these processes. An example of these problems in the Fermilab Tevatron concerns the control of the tune during the various (energy, low-β squeeze, etc.) ramps. These ramps are initiated by global timing signals which start ramps on various magnetic circuits. The quadrupole currents throughout the ramp are set at breakpoints which are separated by several seconds. All control is open loop and the currents at the breakpoints are initially determined from beam measurements. Typically, a satisfactory table describing the current in the circuits as a function of time is established once and is only changed when accelerator conditions change enough so that a "global retuning" is needed.

For reasons which are not well understood (but are the subject of intensive study), the Tevatron operating conditions drift. As a result, a ramp which is optimized one day may not be optimal at another time. The "global retuning" mentioned in the previous paragraph is extremely labor-intensive and cannot be done during ordinary operations but requires dedicated accelerator time. The system we have built (GFSDA for Generic Finite State Data Acquisition) eliminates these problems by providing flexible, large-scale data acquisition which will make correction of tune and other parameters much simpler.

The hardware system we have built initially is VME-based with a Motorola MVME 162-223 processor card with a 68040 processor, a single Omnibyte 4-channel, 12-bit, fast digitizer with onboard memory for 65536 digitizations, and hardware for reading Tevatron clock events. The microprocessor runs VXWorks. Additional ADC cards can be added. The final configuration will be based on a 68060 processor and will have an external shared memory card with 64MB of memory for data storage.

The software system consists of one (or more) Tevatron clock-event-driven Finite State Machines (FSM), a set of states each with its own data acquisition properties, and a set of circular buffers for the storage of the large amounts of data (100s of KB) which the system stores. In addition, the microprocessor will have sufficient computing power to do sophisticated data analysis such as Fast Fourier Transforms (FFT), peak-finding, and closed loop control, if desired.

SOFTWARE IMPLEMENTATION

The Tevatron operations follow the FSM paradigm quite closely. One can envision each ramp in the Tevatron as corresponding to a specific state with its own set of breakpoints for the open loop control,

---

*Operated by the Universities Research Association under contract with the U.S. Department of Energy.

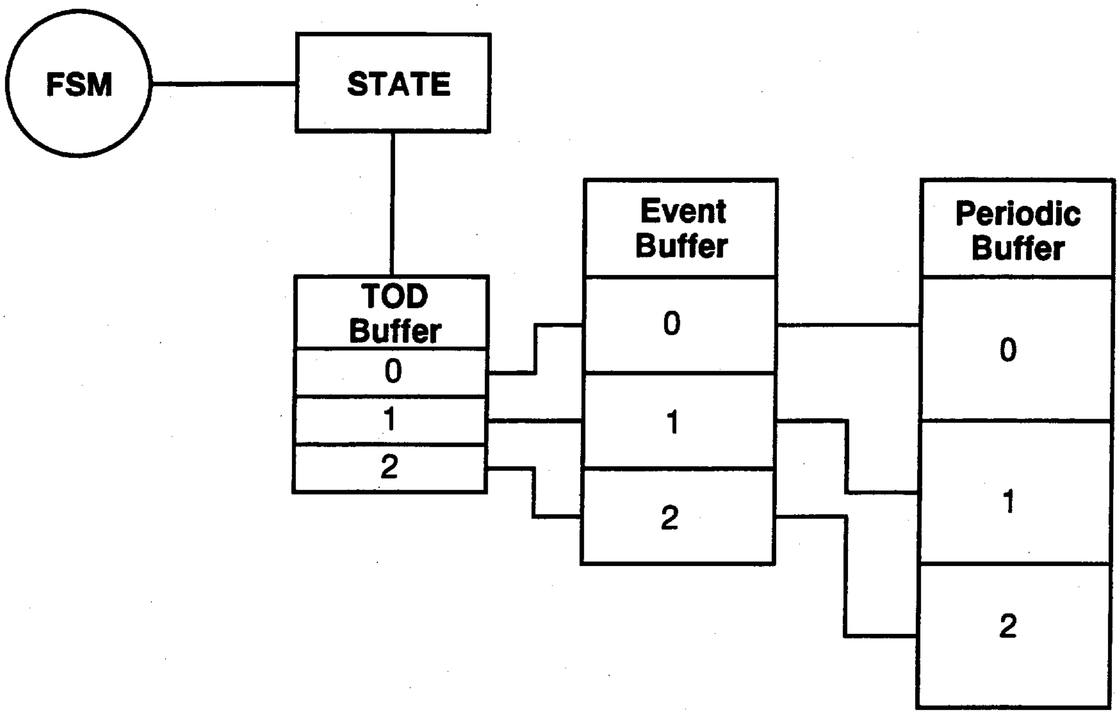

Figure 1. Relationship between the FSM, state, and the various buffers.

its own needs for data acquisition, analysis, and possible closed loop control. The FSM itself simply controls the flow of states, which can be quite complicated as there is not a unique time sequence of operations (for instance, the number of proton injections can vary, certain operations may be omitted, etc.). Transitions between various states are controlled by events (or event + delay) on the global Tevatron clock, by simple time-outs, or by "soft" clock events (a control system parameter). Transitions are allowed to up to 8 different states.

The states control the data acquisition through two types of interrupt rules: event-driven rules and periodic rules. An event rule can be specified as a global Tevatron clock event (or event + delay) or a specified time interval after the start of the state. A given clock event can occur many times during a particular instance of a state. Each event rule has its own specification for data acquisition, including the number of digitizations to be collected and the particular set of ADC channels to collect. Each state can also have a single periodic rule which consists of a frequency for data acquisition as well as the number of digitizations and the ADC channels to be collected.

The data to be collected usually consist of 2 or 4 ADC channels and at least 1024 digitizations. The number of digitizations is determined by the resolution needed for the FFT. The digitization frequency is 62.5KHz (determined by the available internal frequencies for the Omnibyte ADC. Eventually it will be the beam revolution frequency of about 47.713 KHz), so 1024 points gives a frequency resolution of 62 (48) Hz, which is barely adequate for a Tevatron diagnostic. The number of points will be increased once the shared memory is installed. In any case, each event record will be at least 2024 bytes, and possibly much larger. In addition, there can be up to 64 separate events in a given instance of a state, leading to a possible data size for an instance of 131 KB. We also store header information such as the actual event trigger, the time into the state, the ADC channels collected, the number of points collected, etc., in the event record. An intelligent buffering scheme is needed to ensure that these data can be accessed in a simple manner.

The method of data storage we have devised consists of a series of circular buffers. Each state has 3 of these buffers: one each for the time-of-day (TOD) stamps, event-driven data and periodic data. We have designed this system so that the data for many instances of a state can be saved in memory and will be overwritten in a circular fashion. The instances are identified by the TOD stamp, a 28 byte label which gives the date and time of day at which the particular instance was initiated. The periodic and event buffers are assigned at the same time the TOD buffer is assigned. The relationship between the three buffers is such that once a TOD buffer is selected (either for writing by the microprocessor or for accessing from the Fermilab control system), access to the accompanying event and periodic buffers is guaranteed. The relationships are illustrated in Figure 1.

The FSM system is flexible and thus extremely complicated. We have provided a high-level application program running in the Fermilab control system to program the FSM dynamically. It will allow the user to download configurations from a preset menu and create new configurations as the uses of the system evolve. It also enables one to start the FSM and force transitions between states.

Access to the ADC data also occurs through a high-level program which dynamically mirrors the configuration of the FSMs. Currently, the application allows the user to select a particular buffer and do rudimentary FFT analysis and peak-finding.

STATUS OF THE PROJECT

A preliminary version of this system has been completed and is working in the Tevatron as a tune measurement system. The analog inputs are the signals from a set of resonant Schottky detectors. Since these detectors are resonant and have a 1.5μs time constant, they are not really well matched to the digitization and data acquisition capabilities of GFSDA, but nonetheless can be used to measure the tune. This simple system already provides equal accuracy and more flexibility (the ability to capture and store data in real time) than the set of HP 3561A Dynamic Signal Analyzers that have been used for the past 7 years. An illustration of the data we have obtained is shown in Fig. 2. This represents a single digitization of the signals from the vertical Schottky detector.

We intend to make significant improvements to this system. Our initial implementation has a minimum amount of processing in the microprocessor. Data analysis will be done using the Fermilab console system. Once we have developed accurate FFT and peak-finding algorithms we will move them into the microprocessor. This will of course require an expansion of the circular buffer system, but we intend to build upon what we have by maintaining the raw data buffers and adding a set of buffers with FFT and peak data.

In collider operations the Tevatron stores both protons and antiprotons simultaneously. The 2 species see different accelerator lattices and thus can have different tunes, etc. and it is important to measure the tunes separately. The Schottky system in use allows for independent tune measurement with the use of 2 detectors in each plane separated by $\lambda/4$ at the resonant frequency of 21.4 MHz. This separation provides a directionality of at best 20 dB when the analog circuitry is tuned up. In practice, due to drifts, the rejection is closer to 10 dB. GFSDA can be used to record the 4 detectors individually and do the directional rejection in software. This may have several advantages over the analog method now in use: it may be more stable, and it should be possible to do regular (daily) calibrations with only 1 species of particle in the accelerator, ensuring that the rejection is updated.

A more powerful method of measuring proton and antiproton tunes is to use a wideband pickup and gate in time at the instances when the particles pass the pickup. The Tevatron has several stripline detectors and these are being instrumented with sample-and-hold circuits and precise bucket-by-bucket gates which will allow GFSDA to measure the tune of an individual proton or antiproton bunch in the Tevatron.

One of the problems which led to the design and construction of GFSDA was the ill-understood drift in the operating conditions of the Tevatron, leading to changes in tune and other beam parameters. If we can implement a peak-finding algorithm inside the microprocessor, it will be straightforward for us to develop a real-time feedback mechanism through the quadrupole power supplies and control the tune in real time. This has already been done in a simpler prototype system[1].

REFERENCES

[1] D.A. Herrup et. al., Nuclear Instruments and Methods in Physics Research A329 (1993) 470.

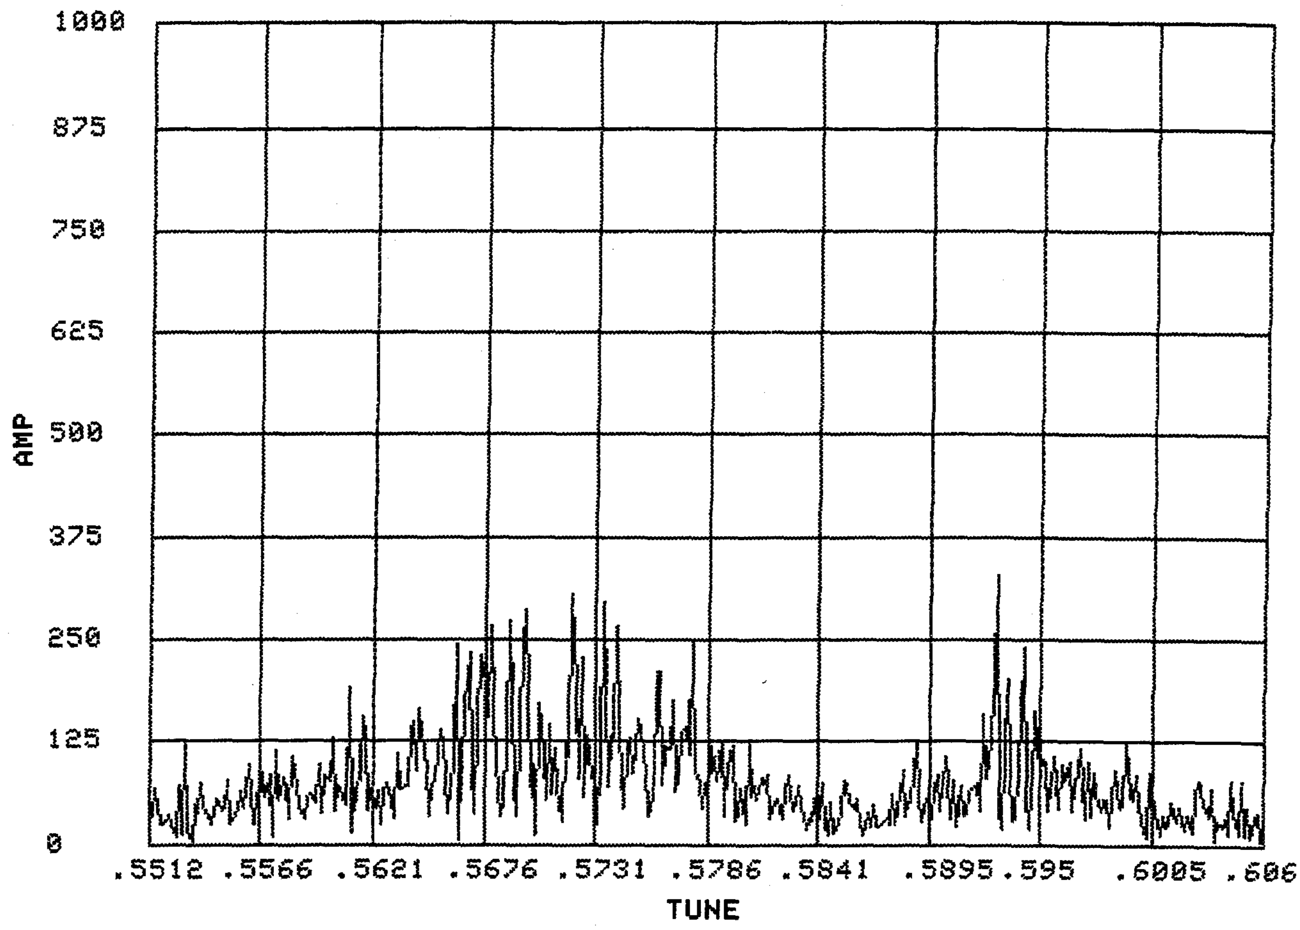

Figure 2. Tune spectrum obtained with GFSDA.